

UNIVERSITY OF CALIFORNIA
Los Angeles

**Cellular Network for Mobile Devices and Applications:
Infrastructure Limitations and Solutions**

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Computer Science

by

Chunyi Peng

2013

© Copyright by
Chunyi Peng
2013

ABSTRACT OF THE DISSERTATION

Cellular Network for Mobile Devices and Applications: Infrastructure Limitations and Solutions

by

Chunyi Peng

Doctor of Philosophy in Computer Science

University of California, Los Angeles, 2013

Professor Songwu Lu, Chair

In recent years, we have seen tremendous growth of mobile devices and applications. The global shipments for smartphones and tablets reached 722M, and 128M, respectively, in 2012. They are predicted to reach 1.52B and 352M by 2017. In the meantime, mobile application downloads in Apple Store and Google Play have already exceeded 50B and 25B, respectively. In order to support such mobile data devices and applications, 3G/4G cellular network infrastructures play a critical role. Indeed, the mobile data traffic carried by the cellular networks are estimated to grow seven times within the next four years, from 1.6 Terabyte (TB) in 2013 to 11.2TB in 2017.

In this dissertation, we study how to improve the 3G/4G network infrastructure to better support mobile devices and applications. Specifically, we focus on two topics: one is how to reduce the operational cost by improving energy efficiency of cellular network infrastructure, and the other is to reexamine the architecture and policy practice on data accounting in 3G/4G infrastructures from both robustness and security perspectives. Both topics complement each other, covering two parts of cellular network infrastructures – radio access network (i.e., base stations) and core network. Given each topic, our methodology is as follows. We identify key problems, use real measurements and traces collected from the operational networks or the mobile devices (if needed) to quantify the impact of the identified issues, analyze their root causes, pro-

pose corresponding solutions, and evaluate their effectiveness. Finally, we share the gained insights and learned lessons.

The first topic seeks to design a green network infrastructure for 3G networks within the 3G standards framework. Our work is motivated by the fact that the 3G cellular infrastructure, particularly the base station (BS) networks, consumes about 80% of overall energy in today's operational networks. Such base station networks incur large energy waste in that their energy consumption is not in proportion to their carried traffic loads. Our study further shows two root causes. On one hand, the traffic volume is not constant over time; In fact, it exhibits high fluctuations both in time and over space. On the other hand, each base station consumes a large portion of energy even at zero traffic, due to its supporting system (e.g., cooling) and idle communication overhead.

In this work, we propose *GreenBSN*, a traffic-driven design to green BS infrastructure. We profile BS traffic, and approximate network-wide energy proportionality using non-load-adaptive BSes. The instrument is to leverage the inherent temporal-spatial traffic dynamics and node deployment heterogeneity, and power off under-utilized BSes under light traffic. We further show that, our solution can be implemented within the current 3G standards. Our evaluation on four regional 3G networks shows that our standard-compliant, practical solution yields up to 53% energy savings in dense large cities and 23% in sparsely deployed regions. The key insight is that, traffic dynamics is pervasive both in time and over space in operational 3G/4G networks, and such diversity opens a new venue for energy savings in large-scale cellular networks.

In the second topic, we examine the data accounting architecture and practice, which has migrated its design from the circuit-switched voice service to the packet-switched data. Our work is driven by the fact that, data-plan subscribers are charged based on the used traffic volume in 3G/4G cellular networks. Though this usage-based charging system seems to receive general success in operation, no effort has been used to investigate it from the user perspective. Therefore, we conduct experiments to critically assess both this usage-based accounting architecture and application-specific

charging policies by operators. Our evaluation compares the network-recorded volume with the delivered traffic at the end device. We have found that, both generally work in common scenarios but may go wrong in the extreme cases: We are charged for what we never get, and we can get what we want for free. In one extreme case, we are charged for at least three hours and 450MB or more data despite receiving no single bit. In another extreme case, we are able to transfer 200MB or any amount we specify for free. The direct root causes lie in lack of both coordination between the charging system and the end device, and prudent policy enforcement by certain operators. The more fundamental problem is that, solutions (e.g., data accounting) that work for circuit-switched, telecom networks may not be directly applicable to the packet-switched IP-based networks. The open-loop data accounting fails to record consistent data volume along the packet-switched path, thus it suffers from various problems in certain worst-case scenarios. We propose remedies that mitigate the negative impacts.

We further study the data accounting problem from the security perspective. We have identified loopholes in both the metered accounting architecture and application-specific charging policies, and discovered two effective attacks exploiting the loopholes. The “toll-free-data-access-attack” enables the attacker to access any data service for free. The “stealth-spam-attack” incurs any large traffic volume to the victim, while the victim may not be even aware of such spam traffic. We also show that, current security mechanisms in cellular networks, such as hardware-based authentication and authorization, and firewalls, cannot defend from such attacks. Our experiments mainly on two operational 3G networks have confirmed the feasibility and simplicity of such attacks. We further propose defense remedies. The main learned lessons are (1) the telecom-based solutions may be inappropriate for mobile data since the virtual circuit does not exist any more. Core network operations should adapt to the underlying PS technology; (2) IP data forwarding uses a push model, causing the data delivery to the mobile victim without its prior consent; Current cellular networks are unable to verify the authenticity of the malicious sender at the network layer.

The dissertation of Chunyi Peng is approved.

Lixia Zhang

Mario Gerla

Yingnian Wu

Songwu Lu, Committee Chair

University of California, Los Angeles

2013

To my parents, who always love me unconditionally

Table of Contents

1	Introduction and Motivation	1
1.1	Popularity of Mobile Devices and Apps	2
1.1.1	Drives on User Demand and Technology Advance	3
1.1.2	Mobile Service Model	5
1.2	Problems in Cellular Network Infrastructure	7
1.2.1	Energy Inefficiency in Radio Access Network	8
1.2.2	Data Accounting Inaccuracy in Core Network	10
1.3	Our Contributions	12
1.3.1	Toward a Green Cellular Base Station Infrastructure	12
1.3.2	Toward Resilient Mobile Data Accounting Architecture	14
1.4	Dissertation Structure	16
2	Background and State-of-the-Art	18
2.1	Cellular Network Architecture	18
2.2	Background on Basic Operations	21
2.2.1	Two Service Models for Voice and Data	21
2.2.2	Handoff for Mobility Support	24
2.2.3	Data Charging/Accounting	25

2.3	State-of-the-art on Cellular Infrastructure	28
2.3.1	Green Cellular Infrastructure	28
2.3.2	Data Accounting	30
2.3.3	Other Aspects	31
2.4	State-of-the-art on Non-Cellular-Network	32
3	Green Cellular Infrastructure	36
3.1	Non Energy-Proportional BS Networks	36
3.1.1	Design Goals and Challenges	38
3.1.2	Roadmap to the Solution	39
3.2	Understanding BS Power Consumption	40
3.3	Characterizing Diversity in Trace Analysis	44
3.3.1	Traffic Diversity Over Time	46
3.3.2	Traffic and Deployment Diversity in Space	49
3.3.3	Capacity and Utilization Diversity	51
3.4	GreenBSN: Towards Green BS Networks	54
3.4.1	Grid-based BS clustering	56
3.4.2	Location-dependent Traffic Profiling	58
3.4.3	Graceful Selection of Active BSes	60
3.4.4	Exploit BS Diversity	63
3.5	Comparing with the Optimal Scheme	65
3.5.1	The Optimal Scheme	66
3.5.2	Bounded performance gap	67
3.5.3	Tradeoffs for practicality	70
3.6	Working within 3G Standards	71

3.7	Evaluation	75
3.7.1	Overall Performance	75
3.7.2	Impact of Various Components	78
3.7.3	Comparing with the Optimization-based Scheme	82
3.7.4	Impact on Clients	85
3.7.5	Evaluation Summary	86
4	Mobile Data Accounting	88
4.1	Problems in Mobile Data Accounting	88
4.1.1	A Motivative Example	88
4.1.2	Issues in Data Accounting	89
4.2	Experimental Methodology	93
4.3	We Get What We Want For Free	95
4.3.1	Loopholes in Charging Policy Practice	96
4.3.2	Toll-Free Data Service Attack	101
4.3.3	Policy as Double-Edged Sword	103
4.3.4	Recommended Quick Fix	104
4.3.5	Progress Update and Carriers in Other Regions	105
4.4	We Pay For What We Do Not Get	105
4.4.1	Extreme Cases	106
4.4.2	Root Cause in Open-Loop Data Accounting	113
4.4.3	Common Cases	115
4.4.4	Recommended Quick Fix	120
4.4.5	Progress Update and Carriers in Other Regions	121
4.5	We Pay For What We Do Not Want	122

4.5.1	Vulnerability Analysis	122
4.5.2	Spam Attack in TCP-based Services	125
4.5.3	Spam Attack in UDP-based Services	127
4.5.4	Root Cause	131
4.5.5	Recommended Quick Fix	132
4.6	Gray Areas in Data Accounting	134
4.6.1	UDP Uplink to a Nonexistent Host	134
4.6.2	Effect of Middle-boxes	135
4.6.3	Packet Drop over the Internet	137
4.6.4	Charging for Application Signaling	138
4.7	Discussion and Conclusion	140
5	Conclusion and Future Work	143
5.1	Summary of Results	143
5.2	Insights and Lessons	146
5.3	Immediate Future Work	150
5.4	Future Plan on Network Support for Mobile	151
	References	153

List of Figures

1.1	Forecast of mobile market growth (2012-2017).	3
1.2	Modules in the mobile service model.	5
1.3	Illustration of mobile service model.	5
2.1	3G/4G Cellular Network Architecture.	19
2.2	3G UMTS charging architecture in PS domain.	25
2.3	Charging process for a data service flow.	26
3.1	Energy-load curves for BS network in Region 1.	37
3.2	A typical base station in 3G networks.	40
3.3	Measurements and model estimates of BS transmission power (P_{tx}) with regards to their carried loads.	41
3.4	Measurement of the cooling power (P_{misc}) at one BS in 2010.	43
3.5	Maps of base station locations in five regions. The right plot in Region 5 plots a 10km x 10km subregion (called as “5A”) where 281 BSes are deployed. Dotted rectangles A and B indicate residential and downtown areas.	45
3.6	Traffic dynamics over time.	47
3.7	Illustration of temporal multiplexing diversity.	49

3.8	Spatial diversity in deployment and traffic.	50
3.9	An example of six-day traffic traces of four BSes in the second data set.	52
3.10	Illustration of BS capacity diversity in Region 5.	53
3.11	Peak and temporal utilization in Region 5 and 5A.	54
3.12	Example of virtual grids. Left: BS on/off status. Right: virtual grids.	57
3.13	Profiling examples.	59
3.14	Example of BS graceful selection.	59
3.15	Illustration of energy saving from dual-RF BS system. The dash lines represents BS coverage and red cells are ON.	65
3.16	Grid-based versus optimization-based schemes.	67
3.17	Workflow of GreenBSN solution.	71
3.18	3G Handoff procedure for user migration.	73
3.19	Evaluation results of various effects on energy-saving, miss traffic, and the number of active BSes.	80
3.20	Comparison with the optimization-based scheme in different cases. Left: case settings; Right: energy saving.	82
3.21	Comparison of our scheme and other designs on energy saving: (a) using real-time traffic; (b) disabling smooth selection.	84
3.22	Transmission range change in Regions 1 and 4.	86
4.1	End-to-end data delivery path and involving components.	92
4.2	Procedure of a typical mobile Web browsing.	97
4.3	Illustration of free Web browsing in the “toll-free-data-attack”.	97
4.4	V_{UE} and V_{OP} in five DNS tests.	98
4.5	Feasibility test of free data services; $V_{OP} = 0$	100

4.6	Three approaches to launch “toll-free-data-access-attack.”	102
4.7	Our indoor testbed, where room A is a dead zone (NS-zone) for both operators.	107
4.8	Procedure of DL-NS experiment.	108
4.9	Accounting volume gaps under various UDP source rates; $t = 1$ minute.	110
4.10	Accounting volume gaps with various in-NS-zone durations.	110
4.11	Illustration of gap creation in various wireless environments.	111
4.12	Gaps under various source rates in different zones in DL-All experiments.	112
4.13	Gaps with with intermittent connectivity in DL-All experiments.	113
4.14	One DL-NS experiment trace using TCP.	116
4.15	Illustration of the stealth spam attack.	123
4.16	Two steps to launch the stealth spam attack.	125
4.17	Wireshark traces at the victim after the UE tear downs the TCP connection.	126
4.18	Data volume caused by TCP-based stealth spam attacks at various source rates.	127
4.19	Data volume caused by TCP-based stealth spam attacks for various durations.	127
4.20	Data volume caused by UDP-based (Skype) stealth spam attacks at various source rates.	129
4.21	Data volume caused by UDP-based (Skype) stealth spam attacks for various durations.	129
4.22	Wireshark traces at the victim after log out from Skype.	130
4.23	Results when connecting to nonexistent hosts using different TCP ports	136
4.24	Results when experiencing unreliable Internet packet delivery.	138

List of Tables

2.1	Table of important abbreviations and acronyms.	19
3.1	Different power models used in this work.	44
3.2	Basic statistics of 5-region traces.	45
3.3	Traffic autocorrelation with 24-hour lag.	48
3.4	Traffic variation in consecutive days.	48
3.5	Max-to-min traffic ratio in neighborhood	50
3.6	Power saving in four regions.	76
3.7	Power saving in peak/idle hours and subregions.	76
3.8	Energy saving with different power models.	78
3.9	The BS-to-client distance change (m).	86
4.1	Example of itemized mobile data usage.	90
4.2	Dimensions of mobile accounting issues	91
4.3	Example results for three DL-All experiments when source rate is $s = 800$ Kbps and $t = 1$ minute.	112
4.4	Volume gaps for applications in DL-NS experiments.	117
4.5	Volume gap for user studies during June 10-23, 2012 (User 7 had only one-day usage record on June 22, 2012).	119

4.6 Signaling overhead of popular applications. 139

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to my academic advisor, Professor Songwu Lu, for his never-ending guidance, support and trust in me throughout the years. Songwu is my role model in many ways. I am very grateful for his high standard for research, care for students, enthusiasm and kindness from the beginning to the end. He created a dream environment to encourage me to explore new ideas freely and pursue high-quality work. He also shared me with his insight, vision and wisdom, while helping me to make a critical decision whenever the project came to a crossroad. Despite his busy schedule, he meets with us regularly, and sometimes on a daily basis when we need his help. Even when I was doing internships, he still took time to talk to me on a weekly basis and provide me many valuable suggestions. He provided amazingly valuable feedback on this dissertation. He will continue to be my role model.

I also have good fortune to work with Professor Lixia Zhang, who has helped me not only in sharing her deep insight in research, but also providing valuable advice in every aspect. I admire her for her broad knowledge, deep understanding and critical thinking in network and systems. Discussion with Prof. Zhang advanced our understanding on architectural issues in cellular networks, and her help directly led to identifying the root cause of the stealthy-spam-attack in Chapter 4.5.

I would like to thank my other committee members: Professors Mario Gerla and Yingnian Wu, for many constructive comments on various parts of my thesis. I also would like to thank Dr. Haiyun Luo, who was my qualifying committee member and have to quit due to his personal reason. He gave me many valuable suggestions and help in the work to green cellular infrastructure in Chapter 3.3.

I also indebted to many mentors during my internships at IBM Thomas J. Watson Research Center and Microsoft Research Redmond. I appreciate the opportunities to work with many excellent researchers: Ranveer Chandra, Kang-Won Lee, Starsky H.Y. Wong, Hui Lei, Minkyong Kim and Zhe Zhang. Ranveer was always helpful with

wireless MAC and rate adaptation issues. He was very friendly to give me advice on research career. I would like to give my special thanks to Kang-Won and Starsky, who provided me with the maximal flexibility, support and guidance in the wireless broadcast rate adaptation project. I also would like to thank Hui, Minkyong, and Zhe, who gave me a fantastic opportunity to work with real problems in product data centers: virtual machine image provisioning. I also appreciate kind help and valuable comments from many other researchers in both labs: Alec Wolman, Ming Zhang, Ramya Raghavendra, Yang Song, Han Chen, Fan Ye, Vugranam Sreedhar and Wei Tan.

It has been a privilege to work together with intelligent and friendly people in the Wireless Networking Group (WiNG) at UCLA. I sincerely thank the group members for their feedback on my work and valuable discussion in various research topics. Many thanks go to Chi-Yu Li, Guan-Hua Tu, Suk-Bok Lee and Yiannis Pefkianakis. I also had pleasure with the company of many UCLA buddies, student interns and old friends. They made my four-year journey enjoyable and memorable.

Finally, I thank my parents for their endless love and support.

VITA

2002	B.E. (Automation), Tsinghua University, China
2004	Inter at Microsoft Research Asia, Beijing, China
2005	M.E. (Automation), Tsinghua University, China
2005 – 2005	Software Engineer on TD-SCDMA physical layer, T3G Ltd., China
2005 – 2009	Associate/Assistant Researcher in Wireless and Networking Group, Microsoft Research Asia, Beijing, China
2010	Visiting student at Peking University (JRI program), China
2011	Intern at IBM T.J. Watson Research Center, NY, USA
2012	M.S. (Computer Science), UCLA.
2012	Intern at IBM T.J. Watson Research Center, NY, USA
2012 – 2013	Intern at Microsoft Research Redmond, WA, USA
2009 - present	Graduate Student Researcher, Computer Science Department, UCLA

PUBLICATIONS

Guanhua Tu, Chunyi Peng, Chiyu Li, Xingyu Ma, Songwu Lu, “Accounting for Roaming Users on Mobile Data Access: Issues and Root Causes”, *To appear at the 11th International Conference on Mobile Systems, Applications and Service (MOBISYS’13)*, Taipei, Taiwan, June 2013.

Ioannis Pefkianakis, Yun Hu, Suk-Bok Lee, Chunyi Peng, Sofia Sakellari, Songwu Lu, “Window-based Rate Adaptation in 802.11n Wireless Networks”, *ACM Mobile Networks and Applications (MONET)*, Volume 18, Issue 1, pp. 156–169, Feb. 2013.

Chunyi Peng, Chiyu Li, Guanhua Tu, Songwu Lu, Lixia Zhang, “Mobile Data Charging: New Attacks and Countermeasures”, *Proceedings of the 19th ACM Conference on Computer and Communications Security (CCS’12)*, Raleigh, NC, USA, Oct. 2012.

Chunyi Peng, Guanhua Tu, Chiyu Li, Songwu Lu, “Can We Pay for What We Get in 3G Data Access?”, *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking (MOBICOM’12)*, Istanbul, Turkey, Aug. 2012.

Chiyu Li, Chunyi Peng, Songwu Lu, Xingbin Wang, “EERA: Energy-Efficient Rate Adaptation for 802.11n Devices”, *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking (MOBICOM’12)*, Istanbul, Turkey, Aug. 2012.

Chunyi Peng, Guobin Shen and Yongguang Zhang, “BeepBeep: A High-Accuracy Acoustic-Based System for Ranging and Localization Using COTS Devices”, *ACM Transactions on Embedded Computing Systems*, Vol. 11, No. 1, March 2012.

Chunyi Peng, Minkyong Kim, Zhe Zhang, Hui Lei, “VDN: Virtual Machine Im-

age Distribution Network for Cloud Data Centers”, *Proceedings of the 31st Annual IEEE International Conference on Computer Communications (INFOCOM’12)*, Orlando, Florida USA, March 2012.

K. R. Jayaram, Chunyi Peng, Zhe Zhang, Minkyong Kim, Han Chen, Hui Lei, “An Empirical Analysis of Similarity in Virtual Machine Images”, *Proceedings of the 12th International Middleware Conference (Middleware’11) (Industry Track)*, Lisboa, Portugal, Dec. 2011.

Chunyi Peng, Suk-Bok Lee, Songwu Lu, Haiyun Luo and Hewu Li, “Traffic-Driven Power Saving in Operational 3G Networks”, *Proceedings of the 17th Annual International Conference on Mobile Computing and Networking (MOBICOM’11)*, Las Vegas, Nevada USA, Sep. 2011.

Alexander Afanasyev, Jiangzhe Wang, Chunyi Peng, Lixia Zhang, “Measuring redundancy level on the Web”, *Proceedings of the 7th Asian Internet Engineering Conference (AINTEC’11)*, Bangkok, Thailand, Nov. 2011.

Jiangzhe Wang, Chunyi Peng, Chiyu Li, Eric Osterweil, Ryuji Wakikawa, Pei-chun Cheng, Lixia Zhang, “Implementing Instant Messaging Using Named Data”, *Proceedings of the 6th Asian Internet Engineering Conference (AINTEC’10)*, Bangkok, Thailand, Nov. 2010.

Suk-Bok Lee, Sai-Wang Tam, Ioannis Pefkianakis, Songwu Lu, Mau-Chung Chang, Chuanxiong Guo, Glenn Reinman, Chunyi Peng, Mishali Naik, Lixia Zhang and Jason Cong, “A Scalable Micro Wireless Interconnect Structure for CMPs”, *Proceedings of the 15th Annual International Conference on Mobile Computing and Networking (MOBICOM’09)*, Beijing, China, Sep. 2009.

CHAPTER 1

Introduction and Motivation

Internet is advancing to a new era of *mobile*. Pervasive mobile data access is coming true for almost anyone, anywhere and anytime, thanks to the globally rapid expansion in advanced mobile cellular networks (e.g., 3G/4G), and the explosive growth of smartphones and tablets (e.g., iPhone, iPad and Android phones). In the meantime, mobile devices and applications are transforming the way how people access and use network services. For example, Facebook and Google admitted the ineffectiveness in their traditional online ad models for mobile users [The12b, The12a].

The popularity of mobile data services calls for cellular network support, which is highly dependable, adaptive, and extensible, for mobile applications. Initially, cellular network infrastructure was designed for voice calls, targeting at the quality of service as the goal of a telecommunication system. To this end, cellular operators invest heavily in the deployment and upgrade of infrastructure equipments (e.g., base stations) to enforce full radio coverage; In order to guarantee the quality of each radio access, cellular networks take a *smart core* approach, performing complex signaling and dedicated control in their closed core infrastructures.

Nowadays, cellular networks are evolving to support mobile services and applications which become the norm. For example, 4G Long-Term Evolution (LTE) networks

have switched to an all-IP based design [3GP11], with the hope to continue the big success of IP for data transmission in the Internet. However, Given the fact that many operations and functionalities still inherit from the conventional telecom-based system, many research problems rise. (1) Do cellular networks well support the increasing growth of mobile applications? (2) Are there any infrastructural limitations, especially those inherent in two different design principles for the telecom network and the Internet? (3) For the current (improper/inefficient) operations, is it possible to improve them without sacrificing the service quality? (4) If yes for the first three questions, what are those limitations and their root cause? How to improve (or even fix) them? What lessons and insights can we learn from the current practice?

In this chapter, we first motivate cellular network support driven by the popularity of mobile devices and applications in Chapter 1.1. Chapter 1.2 describes two concrete design problems in cellular infrastructure. One is the energy-inefficient operation in radio access network, and the other is the improper operation of mobile data accounting in core network. We explore the infrastructural limitations and solutions through these two problems. Chapter 1.3 gives a brief summary of our work. Finally, we provide an organization structure of this dissertation.


1.1 Popularity of Mobile Devices and Apps

The *mobile* era is coming. Recent years have witnessed explosive growth in mobile devices, applications, and usage.

- The global smart connected device (smartphone, tablet, desktop PC, portable PC) crossed 1 billion shipment units in 2012, with 722 million in smartphone, 46.1% more than in 2011, and 128 million in tablets, 78.4% more than in 2011. Looking forward, International Data Corporation (IDC) predicts in 2017, the smart connected device market will continue to surge with 2.2 billion units, with 109.9% growth in smartphones (1.52 billion) and 174.5% growth in tablets (352 million).

On the contrary, the old-fashion desktop PC market is expected to drop by 5% and the portable PC maintains a flat growth of 19% (Source: IDC [Cor13]).

- Apple marked its historic 50 billionth mobile app download in App Store on May 16, 2013 [Bus13], whereas Google announced its 25 billionth download by late September 2012 [goo12]. Nowadays, the revolutionary App Store offers more than 850,000 apps since it was open in July 2008 (with 500 apps then) while Google Play (formerly known as Android Market) provides 700,000 apps.
- The worldwide mobile-to-cellular subscription has exceeded 6.8 billion out of the 7-billion global population in 2012, contributing to 1.5 trillion US dollars revenue for telecommunication services [ITU13]. The year of 2013 will become a breakthrough when the number of subscription first surpasses the population. It brings expansive growth in mobile use. Global mobile data traffic reached 885 exabytes (1EB = 1 million GB) per month at the end of 2012, with 70% annual growth. Moreover, it is expected to increase 13-fold between 2012 and 2017, reaching 11.2 exabytes per month by 2017 [Cis13].



(a) Smart device shipment in [Cor13]

(b) Mobile data volume per month in [Cis13]

Figure 1.1: Forecast of mobile market growth (2012-2017).

1.1.1 Drives on User Demand and Technology Advance

The emerging mobile evolution is driven by both user demand and technology push.

On one hand, users demand “*anytime, anywhere, any device*” services. For instance, people always like to check their emails or surf on the Internet during the occasional waiting time (e.g., waiting for the doctor, having a coffee at the cafe); they even hope to watch videos during a long (boring) journey on a train, subway, bus or even on the flight; People long for a real-time traffic navigation when driving in a crazy-traffic city, say Los Angeles; they also seek a local recommendation of decent restaurants while arriving in a new place; Young people are even keen to post their tweets anytime and quickly follow their friends’ updates (e.g., browsing the newly-uploaded photos); We also expect that each of our favorite settings, contacts, history records, document changes is synchronized and the view is the same from an office workstation or a handheld phone.

Such mobile services possess four key attributes: (a) they move beyond the wired Internet services to ubiquitous computing, such as location-based applications and participatory sensing services; the user context should be taken into account to reshape services; (b) users expect dependable services but without control and management hassle, for example, privacy protection, auto-configurations, software updates, particularly for technology-unsavvy users; Compared with powerful service providers, mobile users look for dumb mobile clients, at least simple user operations; (c) users request for a coherent service platform that delivers integrated functionalities of sensing, communicating and computing; and (d) users can share data and services at will while protecting their privacy in the context of social networking or other online scenarios.

On the other hand, several drivers and enablers are available on the technology push side. They include (1) the wide deployment of high-speed wireless network infrastructures: 3G/4G cellular networks cover 77% of global subscribers [4GA13], advancing to 100Mbps mobile access; in contrast, wireless local area (WLAN) has already evolved from the legacy 802.11a/b/g to 802.11n and beyond, toward 1Gbps+ physical access (e.g., 802.11ac/ad, 60GHz technology); (2) the popularity of fast-evolving portable devices, for example, new generations of smartphones and tablets are released every year with more powerful CPU/GPU, memory and storage, thanks to rapid advance in

multi-core processing and (non-)volatile memory and storage technology; and (3) cloud computing that streamlines control/management over the wired Internet: Ranging from a huge scale of data centers (e.g., 1 million of servers for Google) to medium-sized or small-sized public/private clouds, they impose more efficient resource sharing and utilization while provisioning on-demand services as a utility.

1.1.2 Mobile Service Model

Roughly speaking, a mobile service model consists of three entities: the *front-end* mobile device, the *bridging* network access and the *back-end* services. Figure 1.2 gives a high-level abstraction of the mobile service model, and Figure 1.3 illustrates how to offer mobile applications and services.




Figure 1.2: Modules in the mobile service model.

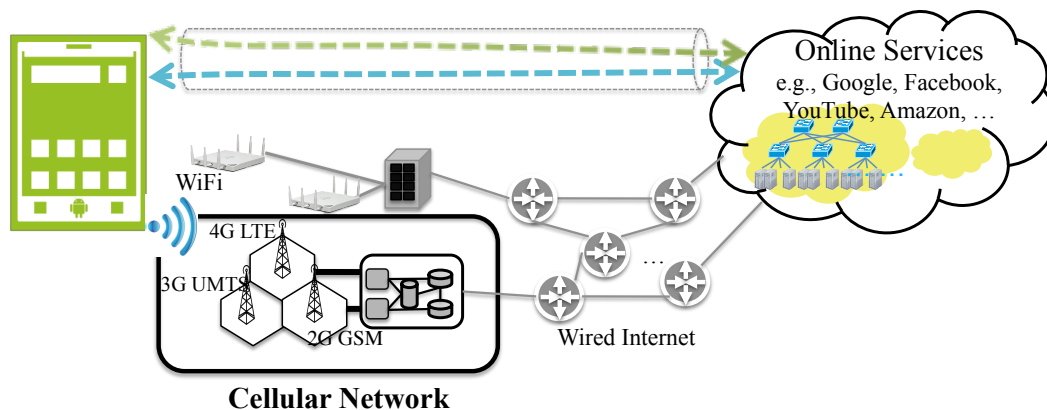


Figure 1.3: Illustration of mobile service model.

The front-end device can be a smartphone, tablet, laptop, or even fancy wearable

devices such as iWatch and Google glass. Heterogeneous devices with various screen sizes, CPU/GPU computation, memory, storage and batteries, take a critical role to offer the interaction to human, rather than conventional desktop PCs. The back-end services provide various online services such as searching, social networking, video-streaming, e-business. These are usually offered via cloud computing, evolving from the conventional server-client practice. The in-between network access (the dotted pipeline in Figure 1.3) offer connections (the arrowed dash lines) between both ends and make all the communication and thus services through.

One revolutionary change from the wired Internet era to the new mobile era, lies in that wireless access to the device becomes the norm. Today, the mobile device uses *only* wireless access (WiFi or cellular networks or both) to indirectly connect with the Internet backbone and finally reaches the wanted services. Compared with the conventional wired access, wireless network infrastructure makes substantial changes to offer highly dependable, adaptive and extensible connections over wireless media. They address a variety of challenges, including limited (much smaller) bandwidth, unreliable connectivity, scalable control and user management, mobility support and so forth.

In this dissertation, we focus on the study of cellular network infrastructure in order to make the mobile devices and applications as first-class citizen. We choose only cellular network infrastructure, not because it is the only challenging component, but because the cellular infrastructure often acts as a major performance bottleneck, among these five modules. Moreover, it stays largely unaddressed compared with a huge amount of research efforts in other areas. Compared with WiFi, an alternative wireless access, cellular network is the only large-scale system to provide universal coverage and mobility support. In the following Chapter 2.4, we will present research problems raised in other four modules (mobile and online services, wired Internet and WiFi), as well as a brief description of research efforts to address these issues.

1.2 ***Problems in Cellular Network Infrastructure***

Cellular network is an indispensable infrastructure to enable pervasive mobile network access. In this dissertation, we aim to explore limitations in the current cellular infrastructure for mobile services. In particular, we would like to address the following questions.

- Does the existing cellular network infrastructure *well support* mobile services, particularly, the rapidly increasing mobile use?
- If no, what are its limitations? And why?
- Given these limitations, is it possible to fix them and how to?
- What insights can we learn, regarding to the design fundamentals of cellular infrastructure?

A cellular network infrastructure consists of two parts: radio access network (RAN) and core network (CN). The RAN provides last-hop radio access to mobile devices. The CN is the central part that provides various services to mobile users who are connected to the RAN. It routes data or voice calls between mobile devices and external networks, e.g., the Internet or the public switched telephone network (PSTN); It also performs critical control/management functionalities, such as mobility support, user authentication, accounting and charging.

Unfortunately, our study shows that both RAN and CN parts face with operational limitations, though they are well designed with good intention. In the radio access network, a large portion (about 80%) of total energy is consumed to power on a huge amount of base stations; such design is used to ensure universal coverage but it turns out that such energy waste might be unnecessary in many scenarios. In the core network, data accounting performed by the operators could be inaccurate; Mobile users might be overcharged or undercharged without bounds in extreme cases. The current accounting

design is intended to provision a simple but effective solution to IP-based data charging, evolving from the traditional telecommunication network. Unfortunately, it turns out that fundamental conflicts lie between the current accounting architecture and IP-based data transmission.

1.2.1 Energy Inefficiency in Radio Access Network

The radio access is realized through cells or base stations (BSes)¹ distributed over land areas. Each BS covers a specific geographic zone under limited wireless coverage (usually several hundred to several thousand meters). Thus, in order to ensure universal coverage, a huge amount of base stations have to be deployed in the spots. As reported in [bs 12, bs 07], the number of global base stations have exceeded 5.9 million by 2012 and doubled in five year.

It raises a growing concern of energy consumption for the cellular network infrastructure [HBB11, WVC12]. Recent reports show that, energy consumption of global cellular networks reached 124.4 TWh in 2011 [ABI08] and is estimated to increase at an annual rate of 15–20% over the upcoming two decades [FZ08]. The energy consumption surge has raised alarms on environmental protection (e.g., increasing CO₂ emission) and operational expense (e.g., huge electricity bills). For example, the power bill for one Chinese mobile operator alone reached \$1.5B in 2009 and is expected to double in five years [Ins10]. The BS subsystem turns to the most critical hinder to green cellular networks. It consumes the dominant portion (more than 80%) of overall energy, while the user clients typically take around 1% [FZ08].

The root cause is that each BS is not energy proportional. More than 50% power is spent on cooling, idle-mode signaling and processing, which are not related to the runtime traffic load. The *always-on* operation causes even larger energy overhead when these BSes are under-utilized, e.g., at mid-night. Our traffic study shows that it is

¹Base stations are interchangeable with cells hereafter.

common that the traffic load at one BS shrinks to 10-20% of the peak one during the idle time (more than 10 hours of a day). In this case, the energy cost irrelevant to traffic loads becomes extremely large, thus resulting in big energy inefficiency.

The *always-on* operation is not designed without rational. Clearly, keeping on all the time can easily guarantee full coverage. Radio signals fade in the air, so each BS has limited coverage and it is able to respond to user requests that are only from its coverage. On the other hand, BS capacity is often configured according to its peak traffic load or its predict. Therefore, it seems that BSes have to stay on in order to offer demanded radio resources when the peak user requests arrive in this area.

However, such *always-on* operation might not be mandatory, at least not at all the time. In fact, peak traffic never comes at all the time. Traffic possesses strongly temporal-spatial diversity and in most time of a day, traffic load is much smaller than the maximal one. It implies that one BS capacity which is configured based on the peak traffic demand, is not always fully utilized. Extra radio resources should be available to serve more user requests under light traffic. Certainly, it requires that the communication to those mobile devices not from its original coverage can go through. On the other hand, though one BS coverage is limited, it is not fixed. During the BS deployment plan, a preferred setting is determined in advance based on the estimate of traffic loads. It happens that each BS needs to control their coverage (transmit power) in order to avoid interfering its neighbors when all the neighbor BSes are on, However, once the neighboring BSes are off or partially off, it is still possible to extend the BS's coverage, especially for those BSes that are densely deployed in the urban hotspots. As a result, BSes can be "*partially on*", which should be able to reduce unnecessary energy waste caused by the *always-on* operation, without the sacrifice of radio accessibility for mobile users.

1.2.2 Data Accounting Inaccuracy in Core Network

In core network, we examine the functionality of data accounting, one basic operation that brings real profits to cellular network operators. We find out that there exists fundamental flaws in the data accounting architecture. Inaccurate data accounting is inevitable, though it may not commonly happen in normal cases or the inaccuracy gap might not be big without intentional (malicious) use.

Most operators apply usage-based (metered) charging for mobile data access. That is, the operator charges the user a monthly bill based on the used data volume. The price for this usage-based charging ranges from 1s to 200s of cents for 1MB data in the US, depending on the chosen data plans. The cheaper the data plan, the higher price the unit data volume. Different from the flat charging scheme over the Internet, 3G/4G operators do not offer unlimited data usage for smartphone users. Both AT&T and Verizon effectively ended such data plans for new customers in 2011, and T-mobile limits the high-speed data volume in its so-called unlimited data plan. Such usage-based charging is not implemented without rationale. The radio spectrum is scarce and mostly licensed, and the offered access speed is bounded by the fundamental limits on channel capacity.

The usage-based accounting is officially stipulated by the 3G/4G standards. Based on the standards, accounting is performed inside the core network on a per-flow basis. Whenever a data flow is initiated with the mobile device, the traffic volume is recorded at the CN when data traverse the CN to reach the mobile or the external server. Therefore, the CN performs accounting operations based on its observed traffic volume.

We assess the mobile data accounting by comparing the user-recorded data volume with the network-recorded usage. Our study is user centric overall, in that *users pay for what they actually get/want at the end systems*. It turns out three categories of inaccurate data accounting occur in reality.

- First, we discover that *we could obtain what we wanted in data access free of charge*. By exploiting the loophole in the enforcement of free-of-charging services, an attacker was able to access any data for free. Note this issue has been fixed since we released our results in 2012.
- Second, we find that *we be charged for what we never receive in certain scenarios*. When packet drops over the last-hop radio link or even the link is totally broken, the accounting gap occurs. In an extreme case when a mobile user roams into a no-signal zone while receiving data from this UDP session without a control loop, the accounting difference of 450MB when this UDP session lasts for three hours!
- Last but not least, we find that *mobile users might be overcharged for those spam data they do not want*. It is easy for an attacker to inject an arbitrarily large volume of spam data into the victim device. No matter if the victim device terminates its data services, the data accounting performed in the core network never stops and the overcharge volume can go unbounded (200 MB observed in our experiment).

These three findings reveal the fundamental loopholes in mobile data accounting. The first is relevant to flow-specific policy that defines what to be charged. Without careful design and implementation, it makes the case to exploit differential charging policy and abuse any transfer that is free or cheaper. The second and third findings are caused by the open-loop data accounting architecture. That is, data volume along an end-to-end path can be inconsistent since no control/signaling is enforced for mobile data access. Without the closed-loop feedback, the accounting system is vulnerable to have different volumes recorded on the core network and the end host. Furthermore, the third finding also reveals another flaw inherent in the accounting system for IP-based mobile data access. IP adopts a push model and delivers any data packets without the receiver's consent. Therefore, the downlink traffic (spam) pushed by an attacker can

easily induce a big mobile data bill to any victim.

We notice that mobile data accounting is among many functionalities offered by the core cellular network. It is still at the early stage to explore limitations in cellular network infrastructure. In this work, we aim to use data accounting as an example to re-examine the cellular core network architecture and gain insights (lessons) to design cellular networks that better support mobile applications.

1.3 Our Contributions

We address the above two problems in cellular infrastructure separately. In the first part to green cellular infrastructure, we explore the root cause of energy inefficiency in cellular base station networks and propose a practical solution to BS energy savings. In the second part, we carry on the first work to investigate mobile data accounting system. We conduct a series of experiments to identify and examine “wrong” accounting scenarios, critically assess the vulnerabilities in data accounting and explore their root causes. Finally, we prototype accounting attacks and propose remedies.

1.3.1 Toward a Green Cellular Base Station Infrastructure

We seek to make cellular base station infrastructure more energy efficient. We first conduct power modeling and traffic analysis using measured BS power consumption, real traffic traces and actual BS deployment maps. These traces are collected from five regional 3G networks from two independent cellular networks, operated by two largest mobile operators in the world. Our analysis reveals that, 3G traffic load exhibits wide-range fluctuations both in time and over space. However, energy consumption of current networks is not load adaptive. The used energy is unproportionally large under light traffic. The root cause is that each BS is not energy proportional. Based on our power model, more than 50% energy are spent on cooling, idle-mode signaling

and processing, in spite of any runtime traffic load. The lighter traffic load, the larger portion of energy overhead.

We thus design GreenBSN, a solution that exploits temporal-spatial traffic dynamics to approximate an energy-proportional (EP) 3G system using non-EP BS components. The main instrument of our proposal is to completely power off under-utilized BSes when their traffic load is light and power them on when the traffic load becomes heavy. The challenge is to devise a practical, distributed solution that uses a small set of active BSes, while satisfying three requirements of traffic capacity, communication coverage, and minimal on/off switching of each BS.

To this end, we take a location-dependent profile-based approach. We divide the network into grids, so that BSes in each local cell can replace each other when serving user clients. We then perform location-dependent profiling to estimate the aggregate traffic among BSes in the grid. Based on the peak/idle of the traffic profile, we decide the corresponding set of active BSes for each duration. It turns out that, if we select the active sets appropriately, we only need to power on a sleep BS and shut down an active BS at most only once during each 24-hour period. We further propose another measure to make use of diversity for energy efficiency. We manually inject diversity by deploying another RF subsystem with different coverage range at one BS. We conduct theoretical analysis and compare our solution with the optimization-based one. It turns out the gap is practically small and it achieves a good tradeoff for simplicity and practicality. Moreover, the implementation of GreenBSN is standard compliant, via reusing cell breathing and handoff techniques which have been enabled at existing BSes.

Our evaluation using real traces shows that our scheme leads to average daily energy saving of 52.7%, 46.6%, 30.8% and 23.4% in the four regional 3G networks (the fifth region is not tested because of limited data sources). The savings are more significant during midnight and weekends and in dense deployment areas, while the miss rate to deny client requests is kept lower than 0.1% in the worst case. While our scheme saves energy on cellular infrastructure, it does negatively increase client power for *uplink*

transmission during *idle* hours (e.g., late nights and weekends).

1.3.2 Toward Resilient Mobile Data Accounting Architecture

We conduct the first work that assesses the mobile data accounting system. We focus on three technical aspects: (1) *What* is charged? It attests the accounting² policy practice by operators; (2) *How* does data access accounting work? It concerns the accounting architecture and its implementation within the 3GPP standards; and (3) Is data accounting *vulnerable* to malicious attacks and how? It examines the accounting architecture from the security perspective.

Our evaluation criterion is *user centric: We pay for what we get (and want)*. We examine the usage gap between the operator-recorded data volume and the user-logged data amount. We conduct experiments with smartphones on two major US operational 3G networks, while also running similar tests in the third US carrier and two carriers in China and Taiwan. Using the phone-logged traces and the data volume recorded by the 3G accounting system, we analyze accounting behaviors in various scenarios.

Our study validates all the three problems described in Chapter 1.2.2. First, we discover that *free access for any data was possible*. It exploited the loophole in the free-of-charge service (e.g., DNS) policy by all three US operators. By constructing a “DNS tunnel” for other data transfer, we were able to transmit 200MB or any amount we specify for free. Based on this, we built the prototypes of the *toll-free data access attack*. The root cause is that, operators use application-specific charging policy, and loopholes exist in such policy practice. Operators do not enforce the full flow-based charging scheme by the 3GPP standard, but using only one or two fields in the five-tuple flow ID. Our work also shows that the policy enforcement is indeed operator specific. While all three US operators offer free DNS, the Chinese operator and the Taiwan carrier still charge it as usual. We proposed the immediate fix to stop free DNS

²We do not differentiate accounting from charging in this work by a slight abuse of wording definition.

policy, followed by these three US operators.

Second, we observe that *we might be charged for what we never receive when the last-hop wireless link fails*. The difference is generally small in the normal cases, resulting in about 10s to 100s of KB in typical applications. However, it can go up to 10s of MB for certain applications (e.g., video streaming). Moreover, the gap can grow quite large in extreme conditions. In one extreme case with a UDP session in a no-signal zone, 450MB is charged despite no single bit ever received by the user. We also identify its root cause. It turns out that, current cellular accounting standards do not explicitly take feedback from end devices. Accounting action is taken at the core components (e.g., GGSN/SGSN in 3G UMTS) inside the core infrastructure. The core components simply record the data volume traversing them to/from the given user for accounting purpose. Consequently, whenever packet drops occur after traversing these components in the no-signal scenario, the accounting system does not know the device status and overcharging may arise. From the user perspective, other cases also exist. For example, extra charge occurs for lost packets due to unreliable wireless channels, or the fake end-to-end connection to the HTTP/FTP proxy deployed by one operator, or sending UDP data to an Internet host with an invalid IP address. The solution fix is to take feedback directly from the end device or access the device status information already collected within the infrastructure (e.g., at RNC) when making accounting decisions.

Third, we identify one new type of attack, *stealth spam attack*, against the accounting system. In the attack, malicious users will inject arbitrarily large volume of spam data into the victim device, even after the target device has terminated its data service, thus fully unaware of such a spam. This attack exploits the architecture weakness of not using feedback from the user when making charging decisions, as well as features of Internet instant messaging applications (e.g., Skype and Google Talk). The fundamental root cause is that IP-data forwarding uses a push model and makes it easy bypass the authentication and security mechanism to protect mobile user accounting.

Our prototypes show that such an attack is feasible and simple enough to launch over the operational cellular networks. Our experiments on two US cellular networks show that, the overcharge traffic volume can go unbounded.

Four main contributions of this work are as follows: (1) We report the first analysis on the 3G/4G network accounting system and identify its loopholes; (2) We identify undercharging and overcharging cases and use real experiments over operational 3G networks to investigate their damage in extreme cases and normal use; (3) We describe two new types of attacks, i.e., *toll-free-data-attack* and *stealth spam attack*, which exploit the identified loopholes to undermine the accounting system; we also use real experiments to validate the feasibility and simplicity of these attacks and their potential damage; (4) We articulate the root cause for the existence of these loopholes and propose effective solutions to eliminating them.

1.4 Dissertation Structure

Now we lay out the structure of this dissertation. Chapter 2 introduces background of cellular network architecture and relevant operations. We then summarize the related work to green cellular network, mobile data accounting, as well as other aspects in cellular infrastructure. In particular, we present two approaches to green cellular network: improved component technology and dynamic cell management. Regarding mobile data accounting, we come up with the first research work and thus we introduce some follow-up work and other mobile research problems related to mobile security and architectural issues. At the end, we give a brief summary of state-of-the-art in non-cellular-network modules.

Chapter 3 presents our efforts to green cellular infrastructure. We first present the problem of energy inefficiency in base station networks: the energy consumption is not proportional to its carried traffic load. We then model power consumption at a base station based on real measurement; It demonstrates that single BS is non-energy-

proportional due to a large power overhead at the cooling and basic communication processing that are irrelevant to traffic loads. Later, we analyze 3G traffic and deployment using two independent cellular network trace datasets. We demonstrate that cellular networks yield a high temporal-spatial diversity, which imply design insights to the green solution. Following these measurements, we describe the proposed GreenBSN solution, its theoretical analysis compared with the optimal one, as well as its implementation within the 3G standard. Finally, we evaluate the GreenBSN performance based on real-trace driven simulation.

Chapter 4 presents our efforts to assess mobile data accounting. We start with an example of mobile accounting and then describe the problem statement and study methodology. We then report four findings in mobile data accounting. We first report the extreme undercharging cases where users might be never charged for data they have delivered, and then demonstrate how to exploit the loopholes to build toll-free data services as an attack against mobile operators. Note that, this undercharging work is not valid since the operators have fixed their loopholes when we released our finding. We then move to another extreme cases where users are charged for data they never receive. We explore the root cause inherent in the open data accounting architecture and propose the remedies. Following this overcharging case, we describe another attack that arbitrarily increases mobile data usage of the victim by stealthily injecting spam packets. We also explore the root cause and propose the fix solutions. The last finding describes other gray-area cases in mobile data accounting. Finally, we further discusses architectural and policy issues.

Chapter 5 conclude with our contribution and directions in future work. Our contributions not only lie in two concrete study to address energy efficiency and data accounting in cellular infrastructure, but also in the exploration of the cellular infrastructure fundamentals. We share the lessons and insights we have learnt. After that, we present the immediate future work and the outlook of network research for mobile.

CHAPTER 2

Background and State-of-the-Art

In this chapter, we first introduce the background of cellular network infrastructure, including network architecture and basic operations involved in this dissertation. Then, we present the state-of-the-art of cellular network infrastructure research. We mainly describe related work in energy efficiency and data accounting, as well as other aspects in cellular infrastructure. Finally, we briefly summarize research progress in other components (mobile, online service, Internet, WiFi), in addition to cellular networks.

2.1 Cellular Network Architecture

Cellular network has two main parts of radio access network (RAN) and core network (CN), as shown in Figure 2.1. In this work, we study the most popular 3G/4G systems, that is, 3G Universal Mobile Telecommunications System (UMTS) and 4G Long-Term Evolution (LTE). UMTS is most widely deployed 3G cellular network technology that offers both data and voice services [Ame10] while LTE is the only mainstream 4G standard. Note that, the overall architecture is similar (if not the same) in all types of cellular networks. We use 3G UMTS/4G LTE as the context to introduce network equipments and their functions; the idea of our work can be applied to other cellular

networks without loss of generality. For reference, Table 2.1 lists important acronyms used in this dissertation.

BD	Billing Domain	GPRS	General Packet Radio Service
BS	Base Station	PDP	Packet Data Protocol
CDR	Charging Data Record	PS	Packet-Switched
CN	Core Network	RAN	Radio Access Network
CS	Circuit-Switched	RNC	Radio Network Controller
EH	External Host	SGSN	Serving GPRS Support Node
FBC	Flow Based Charging	UE	User Equipment
GGSN	Gateway GPRS Support Node		

Table 2.1: Table of important abbreviations and acronyms.

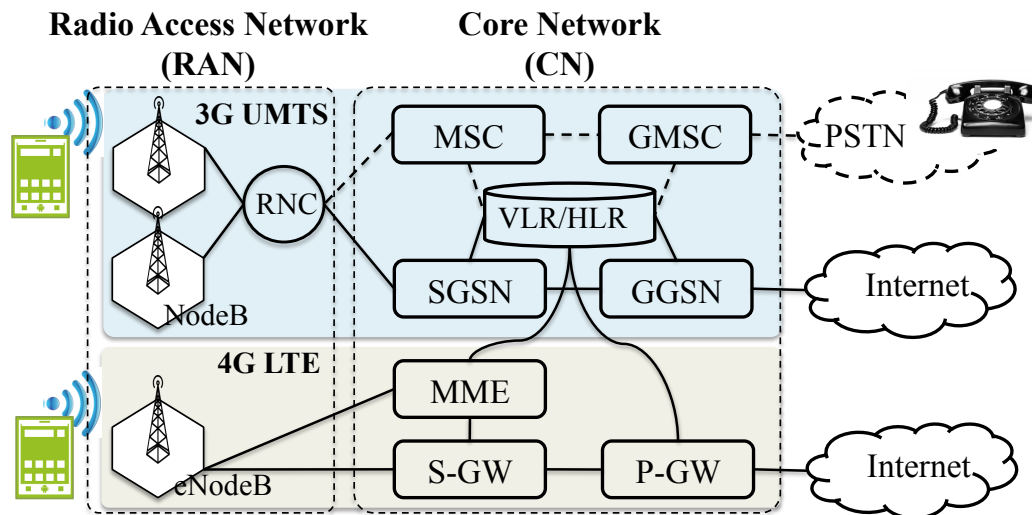


Figure 2.1: 3G/4G Cellular Network Architecture.

3G RAN is composed of the User Equipment (UE), the Base Station (BS)¹, and the Radio Network Controller (RNC). Each RNC manages tens or even hundreds of BSes via Iu-B interface, while exchanging calls/sessions and provisioning services with the core network. It provides central control for radio resource management in RAN,

¹It is also called NodeB in the 3G context and eNodeB in 4G context.

including radio resource control, admission control, channel allocation, and mobility management. Each RNC communicates with the Packet-Switched (PS) and the Circuit-Switched (CS) core networks to provide data and voice services. The BS is the physical unit providing network access services to mobile users via its air interface to the UE. The main functionalities of a BS include wireless link transmission/reception, modulation/demodulation, fast link adaptation, packet scheduling, physical channel coding, error handling, and power control.

3G CN support both CS and PS for voice and data, respectively. The major components of the CS core network are Mobile Switching Center (MSC) and Gateway Mobile Switching Center (GMSC). The former is responsible for paging the UEs in the CS domain and establishing voice calls, whereas the latter routes voice calls between the MSC and the PSTN. The major components of the PS core network are the Serving GPRS Support Node (SGSN) and Gateway GPRS Support Node (GGSN).

SGSN is responsible for the delivery of data packets from and to the UEs within its geographical service area. SGSN records the relevant location information (e.g., current cell) and user profiles (e.g., IMSI, addresses used in the packet data network) of all mobile users registered with this SGSN and performs mobility management, logical link management, authentication, encryption, compression and charging functions. GGSN serves as the hub between the SGSN and the external data networks, e.g., the wired Internet. The GGSN is responsible for IP address assignment and performs functions similar to a router for the connected UE. From an external network's point of view, the GGSN is a router to a sub-network, because the GGSN 'hides' the 3G UMTS infrastructure from the external network. The GGSN also performs authentication, IP Pool management, QoS enforcement and charging functions. The GGSN is connected with SGSNs via an IP-based backbone network. There are two important components critical to user operation and management. Home Location Register (HLR) and Visitor Location Register (VLR), retain user profiles and location information for mobility management, user authentication, and subscriber-specific charging based on types of

mobile subscriber.

LTE is a 4G cellular network standard, offering even higher speed. Its architecture is similar to 3G UMTS. The major difference (also shown in Figure 2.1) is that it only supports PS. Thus, the functionalities of SGSN and GGSN are performed by Mobility Management Entity (MME), Serving Gateway (S-GW), and Packet Data Network Gateway (P-GW) [HT11]. S-GW/P-GW are responsible for routing data packets from RAN to the external data network, whereas MME manages user mobility, e.g., tracking and paging the UEs). In RAN, the functionalities of RAN are performed by eNodeBs (not decoupled at NodeB and RNC).

2.2 Background on Basic Operations

We next briefly introduce major operations or functions involved in this dissertation. They include Circuit-Switched and Packet-Switched technologies that enable two services models for voice and data, handoff for mobility support, and basic process of mobile data charging.

2.2.1 Two Service Models for Voice and Data

Both voice and data are killer applications for today's mobile users. For decades, cellular networks have been using circuit-switched and packet-switched technologies to provide two different services models for them.

To ensure guaranteed service quality, voice has been traditionally via the circuit-switched technology, which establishes a virtual circuit and reserves resources for each call. In contrast, data has been delivered by the packet-switched technology over the Internet, which offers best-effort service without throughput and delay assurance. In the cellular network context, 3G networks have been using the dual-mode infrastructure, with both CS and PS in operation to support voice and data, respectively. Current 4G

LTE networks have adopted an PS only, all-IP based system [3GP11], which works well with the increasingly popular mobile data applications and traffic. For voice support, it has two major solutions of Circuit-Switched Fallback (CSFB) [3GP12a] and Voice over LTE (VoLTE) [VoL]. CSFB leverages the deployed, legacy 2G/3G network that uses CS. Whenever a voice call request is made, the call will be relayed to the 2G/3G network and supported via its CS core. VoLTE is purely PS-based and offers voice service via VoIP.

Different from the Internet, cellular network evolves along a distinct path with different design goals, principles and approaches. Originally in the 1970s, cellular network was designed only for the purpose of telecommunication. The first generation (or 1G) was based on analog system transmission and slowly grew until the 1990s when 2G was introduced. 2G deploys Global System for Mobile communication (GSM) and uses digital modulation to improve voice quality and offer limited data service. Afterwards, cellular network has been following the digital direction, from 2G to 4G, while offering various services, such as paging, faxes, voicemail, text messages and finally data services. Though data becomes increasingly important, CS has been widely used as one dominant technology for a long time, ruling and shaping the cellular network in many aspects.

To guarantee service quality, CS networks reserve a dedicated point-to-point channel for the entire communication. CS signals pass through all the switches (gateways) to reserve resource before a call connection is established. In 3G context, these devices are the UE, the BS, the RNC, the MSC, the GMSC in turn. During the voice call, no other traffic is allowed to use this reserved channel. To this end, many types of connections are established, including Radio Link (RL), Radio Bearer (RB), Radio Access Bearer (RAB), and Radio Resource Control (RRC). The RL is the physical connection between the UE and the BS, while the RB represents the logical connection between the UE and the RNC. The RAB represents the logical connection between a UE and the core network (i.e., MSC for CS call, SGSN for PS data). RAB is service specific;

a UE that is simultaneously using multiple services can have multiple RABs in an on-demand fashion. The RRC connection is used to signal between a UE and an RNC, and it is always first created before establishing any actual voice (or data) session. It is used to control the state and functions of radio resources. One single RRC manages multiple RAB connections. With dedicated control and signaling, states at all the network equipments are well controlled and synchronized for voice calls. This also makes the case for a variety of control and management functions in the core cellular network, such as authentication, mobility management, QoS control.

CS is designed for quality-guaranteed voice calls but it does not fit well with data transmission. Full occupation of resources results in a big waste for data delivery. Therefore, PS applies statistical multiplexing to utilize network resources as possible [CK74]. That is, messages are broken into small, separate blocks (packets) that seek out the most efficient (different) routes based on the destination address in each packet; Network resources are independently managed by statistical multiplexing at each hop. Such idea became the foundation of the Internet. Cellular network adopts the PS idea to support mobile data, while still retaining old features. For example, to start a data session, it needs to establish the above connections (RL, RB, RAB, RRC) in advance. Moreover, Packet Data Protocol (PDP) Context² should be activated before any data connection [3GP06a]. During the data connection, RRC is performed to control radio resources all along. Many underlying control signals run to make the connection through cellular network.

In a nutshell, cellular network offers two distinct service models together, the quality-guaranteed one and best-effort one. The co-existence of current CS and PS design inherits nice features (hopefully) of traditional cellular network as a successful telecommunication system and the Internet, the prevailing data networks. It makes it possible for cellular networks to quickly shift from serving dominant voice calls to accommo-

²PDP contexts provide all the required information for IP packet data connections in cellular networks.

dating the increasing data demands. However, it may also impose risks once these two designs conflict with each other. We will address it in the example of data accounting in Chapter 4 and discuss its insight in Chapter 5.

2.2.2 Handoff for Mobility Support

Cellular network is the only large-scale network system that provides universal coverage and mobility support. Handoff is an indispensable, enabling feature.

The transmission range of one BS is limited (e.g., ranging from several hundred of meters to several kilometers). To provide seamless mobile access, the cellular operator deploys a large number of BSes, each covering a small area. When a mobile device roams from the coverage of one BS to another, it performs a handoff to switch its associated, serving BS. The handoff procedure works as follows. It is first triggered by the mobile device or the serving BS when needed (e.g., when the perceived signal strength is too weak); then, the serving BS (say, BS1) finds another BS (say, BS2) that probably provides better performance (e.g., with stronger signal strength); it sends BS2 a handoff request to reserve radio resources for the mobile device. After that, the mobile device disconnects from BS1 and connects to BS2. During this process, internal gateways also update mobile location and adjust the forward path accordingly. The ultimate goal of a handoff is not to interrupt ongoing services.

Though its concept is relatively straightforward, the implementation is quite complicated in reality. With various cellular technologies and BS types involved, there are a number of different types of handoffs, for example, inter-cell and intra-cell handoffs, depending if the handoff happens within the same BS. They also can be divided into hard and soft handoffs. A hard handover is one in which the channel in the origin BS is released before the channel in the target BS is engaged. A soft handoff is one in which the channel in the origin BS is retained and used for a while in parallel with the channel in the target BS. In this case, two connections are used in parallel, though the interval

may be brief or substantial.

Handoff is one crucial technique for mobility support. In this work, we will reuse its procedure for the purpose of energy saving.

2.2.3 Data Charging/Accounting

Data charging is a crucial function performed by cellular operators. We now introduce the charging architecture and process for mobile data services. Unlike the flat charging practice over the Internet, the current cellular network has been using usage-based charging for its data services. That is to say, the operator collects the actual usage volume over time for each user (i.e., accounting) and imposes charges accordingly.

Broadly speaking, accounting is performed on a per-connection basis. To communicate with a host on the Internet, the mobile device needs to first create a bearer service connection with the cellular network, which is further connected with the wired Internet. Once the connection is established, data packets are delivered. The connection has to traverse gateway-like devices (similar to routers in the Internet) in the cellular network core. These gateways then perform the accounting operations by recording the data volume of those packets that traverse them, until the connection is completed.

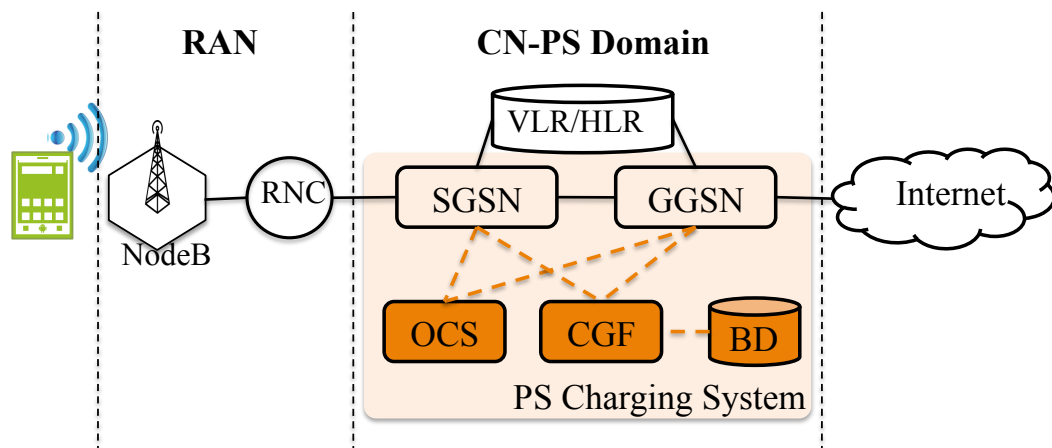


Figure 2.2: 3G UMTS charging architecture in PS domain.

Here, we mainly use 3G UMTS to describe how charging actions proceed in the core network. Note that, the mechanisms are also applicable in 4G Long Term Evolution (LTE) with S-GW/P-GW replacing SGSN/GGSN. In 3G UMTS, SGSN and GGSN are these gateways that perform accounting operation while data packet traverse them. In addition to SGSN and GGSN, three more charging components work in the PS domain: the Billing Domain (BD), the Charging Gateway Function (CGF), and the Online Charging System (OCS), as shown in Figure 2.2. It offers two charging methods in offline and online modes [3GP06c]. In offline charging, data usage is collected during service provisioning in the form of Charging Data Records (CDRs) [3GP07], which are sent to the BD to generate data bills offline. SGSN and GGSN are responsible for collecting data usage and generating CDRs. CGF is used to validate CDRs from SGSNs/GGSNs and transfer CDRs to the BD. This is the most common charging way since mobile users can proceed their data services without limitation. In online charging, mobile users have to pre-pay to obtain credits for data services in advance. Different from the offline charging, user credits must be reserved before data delivery. The OCS authorizes whether or not users have enough credits. GGSN/SGSN deducts data usage from the available credits and stops data services upon zero credit.

We next describe how mobile users are charged for data services, through an example given in Figure 2.3. Consider offline charging, and Alice is about to upload one photo to her Facebook, thus starting a PS service (say, HTTP).

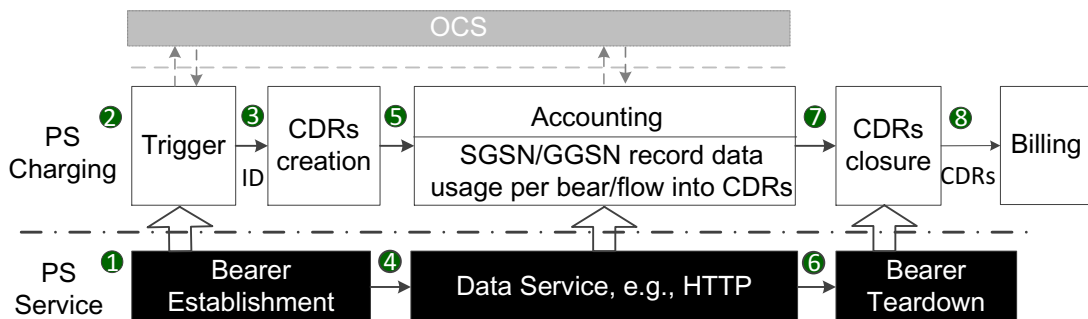


Figure 2.3: Charging process for a data service flow.

Initially, Alice has no available bearer service connection (which may carry one or multiple PS services). She thus establishes a bearer via Packet Data Protocol (PDP) Context³ Activation [3GP06a] (Step 1). Upon this activation, the UE device is allowed to connect with the external data network through the SGSN and GGSN. This activation also triggers the charging procedure, and GGSN assigns a unique charging ID to the activated PDP context (Step 2). SGSN and GGSN then start to create CDRs using the charging ID (Step 3), and are ready to record the upcoming data volume. In addition to charging per PDP context, 3G also supports charging per data flow, called as Flow Based Charging (FBC). FBC separates charging for different services (e.g., web or VoIP) within the same PDP context [3GP08a]. One data flow is typically identified by the five-tuple: (1) source IP address or mask, (2) source port number, (3) destination address or mask, (4) destination port number, and (5) protocol ID of the protocol above IP, e.g., TCP or UDP [3GP06b]. For example, a HTTP data flow can be represented by (*, *, *, 80, TCP). Each of the five tuples can be a wildcard.

With the established PDP context, Alice can upload her photo to Facebook. Both SGSN and GGSN route the UE's packets to/from the external data network during the data service session (Step 4). In the meantime, SGSN and GGSN record the traffic volume that arrives at them into corresponding CDRs (Step 5). Both SGSN and GGSN count the payload of GTP-U (GPRS Tunneling Protocol- User Plane) packets as data volume; GTP-U delivers data within cellular networks and runs below the IP protocol. Therefore, the data volume counts all packet headers above IP, including IP, TCP, and HTTP headers, but the MAC header is not counted.

The accounting procedure (Step 5) lasts until this data service completes. It occurs when the UE tears down this bearer (Step 6) in bearer-based charging, or when Alice closes her HTTP session in flow-based charging. CDRs⁴ are subsequently closed and

³PDP contexts provide all the required information for IP packet data connections in cellular networks.

⁴For the sake of efficiency, CDRs are maintained in the volatile storage (e.g., memory) in practice. Thus, when the CDR size is beyond one threshold, it may be closed and transferred to the BD directly and a new CDR associated with the same charging ID is created. Finally, the BD merges multiple CDRs.

transferred to the BD (Step 7). Finally, BD generates a billing item for the proper user based on the charging ID.

The online charging process is similar, though OCS participates in the triggering and accounting steps (Steps 2 and 5) by authenticating the GGSN/SGSN to use user credits. There is also no need to send CDRs to generate a bill since the consumed credits have been deducted (see Figure 2.3). In the work, we focus on the offline charging, and the same issues also arise for the online charging.

4G LTE applies similar charging architecture and process except the SGSN and GGSN are replaced by S-GW and P-GW, with similar functions. Moreover, bearer establishment is supported by two procedures, i.e., Evolved Packet System Bearer Activation and Public Data Network Connectivity Procedure [3GP11]. The charging system for LTE is almost identical to 3G UMTS, with S-GW and P-GW (replacing SGSN and GGSN) in charge of collecting data usage and generating CDRs.

2.3 State-of-the-art on Cellular Infrastructure

2.3.1 Green Cellular Infrastructure

Energy efficiency in the cellular infrastructure has been an active topic [HBB11, WVC12, Hua, Eri08, EL04, Pan12, Ins10, MCC09, DBM10, BNC10, SKY11, FBG11, OK10, RFF09, BOL09, NWG10, OKL11]. [HBB11] and [WVC12] give a good tutorial on green cellular networks by summarizing existing techniques and addressing research challenges. Generally speaking, proposed solutions can be classified into two categories: *improved component technology* and *dynamic cell management*.

The first approach is to improve energy efficiency of various BS components, including BBUs with standby mode [Eri08], optimized cooling [EL04], more efficient power amplifier [Hua] and green energy (solar or wind energy) supply [Pan12]. These solutions focus on individual component technology and do not aim to explore system-

wide energy saving. They also complement our scheme. C-RAN [Ins10] proposes to use a cloud-based architecture for energy savings. It deploys RRUs in the field but aggregates all the BBUs into a data center, which uses centralized cooling and traffic management to reduce energy consumption. However, C-RAN requires an overhaul of current 3G infrastructure. It is not compatible with the current cellular architecture and standards.

The second approach is to utilize dynamic cell management, e.g., adjusting cell size while turning off idle BSes. Existing studies mostly focus on the theoretical side by seeking to solve various forms of optimizations [MCC09, DBM10, BNC10, SKY11, FBG11, OK10] or describes one specific energy-saving technique [RFF09, BOL09, NWG10, OKL11]. Specifically, [DBM10] formulates the optimal user-BS association as a binary integer programming problem. [MCC09] studies the optimal time to power off BSes by assuming that all BSes power off simultaneously. [BNC10] studies the cell-size optimization given the traffic load. [SKY11] addresses a cost minimization problem that trades-off between energy efficiency and flow performance, and [FBG11] formulates it as a power assignment problem for a specific time interval. [OK10] analysis the energy saving with regards to the on/off time switch. All of [RFF09, RFM10, BOL09] discuss the impact of cell sizes on energy saving and propose a green BS deployment. [NWG10] introduces cell zooming (i.e., cell breathing) to adjust cell size to save energy. [OKL11] uses a case study to show that dynamically switching on/off is able to save energy. Conceptually, our work belongs to this category but aims to build a *practical* solution that is able to work in the operational cellular networks. We differ from them in three aspects. First, we address the problem of energy saving based on real traces and measurements taken from operational networks, without making idealized or simplistic assumptions. Second, we devise a novel, grid-based profiling approach, which is distributed rather than centralized. It is simple and effective in practice. Third, we identify and assess various practical factors ignored by early studies in the power-saving operations. Another practical, follow-up work is to offer

virtual coverage on demand, rather than fixed coverage, to save power [HAB13].

2.3.2 Data Accounting

Despite the popularity of 3G/4G data services, mobile data accounting remains a largely unaddressed area in the research community. [Ezz05] offers a nice tutorial on pricing, charging, billing methods for 3G systems up to 2005. [SJH12] offers recent survey on pricing models and their another work [HSJ12] proposes Tube, a time-dependent pricing scheme to defer mobile data transmission during heavy congestion; Both are orthogonal to the accounting issue studied in this work. To our best knowledge, we conduct the first piece of research on mobile data accounting. We study various cases of overcharging and undercharging in 3G networks in both normal, extreme and attack scenarios. Following our work, [GKW13] studies accounting accuracy for TCP retransmission.

Among current industry efforts, Cisco proposed overcharging protection for GGSN to be aware of lost coverage at end devices based on the feedback of SGSN [cis]. Before our study, another way to obtain free data service in certain carriers is reported to use certain, free Access Point Name (APN), e.g., AirTel India [frec] and UK Three [free]. Our approach to obtain free data services is through a DNS tunneling technique, which is also used in the iodine tool [Iod], dns2tcp and NSTX [ip]. They are designed for data access in scenarios different from ours, where DNS queries are allowed but the Internet access is blocked. They are similar to our toll-free-data service approaches in principle, and we show such ideas also work in wireless cellular networks.

In the more general context, Internet accounting and pricing have been explored in the literature [KP02, TB10, CSE93, SSO08] (see [Odl01] for a survey for work up to 2001). [KP02] provides a taxonomy on Internet accounting and proposes to collect cost objects based on duration, volume, distance, or network resources. [TB10] discusses the pricing architecture for the future Internet where we should pay for information,

services, content, rather than connectivity. [CSE93] and [SSO08] explore the theoretical side on how pricing affects network performance and revenue. These prior efforts focus on the wired Internet. The proposed accounting solutions are quite different from the one used by current cellular networks.

2.3.3 Other Aspects

We briefly summarize other research topics in cellular infrastructure.

Security and privacy have been important and active areas in cellular infrastructure. [ETM05, TEM06] demonstrate how to overload the control channels of 2G/3G networks, thus blocking the legitimate Call/SMS service in the target area. [TML07, RHS06, LBW09] study how to block the CS service caused by the unwanted traffic in the PS domain. [TLO09] measures the impact of malicious mobile devices on the core cellular network. [QM12, QMX12] show that current cellular infrastructures expose security loopholes of TCP hijacking due to their NAT/firewall settings. [KKH12, QWX12, SSL13] show that cellular infrastructure might leak user privacy information (e.g., location) and even further hurt user performance. [AMR12] verifies loopholes of 3G protocols that leak user identification. Moreover, [ZB11] demonstrates that user privacy cannot be achieved via anonymization (Daily activities reveal user identifications), according to a large scale of data mining.

Energy efficiency of cellular network operations has also attracted a lot of attention. Since the energy tail was identified in the RRC control cycle [BBV09, QWG10b], a large number of follows-up have been proposed to reduce energy consumption at mobile by eliminating tail energy. For example, [QWG10a] proposes Tail Optimization Protocol (TOP) to reduce energy tail whenever possible by enabling cooperation between the mobile phone and the RAN; [ABG12] design and implement RadioJockey, a system that uses program execution traces to for the tradeoff of energy saving and signaling overhead in fast dormancy (forcing the mobile radio to quickly go into a low

energy state after a fixed short idle period). [SNR10] uses opportunistic scheduling, which exploits stable wireless channels under regular user activities, to avoid high energy consumption under weak radio signals. [SNR09] uses WiFi offloading to reduce cellular energy consumption. [HQG12] examines power consumption of 4G LTE networks

Recent study in cellular network infrastructure also address other problems, such as network architecture and operation optimization. [BCJ12] proposes CloudIQ, a framework for processing base stations in a data center, similar to C-RAN [Ins10]. [WQX11] conducts the study of middlebox (NAT boxes and firewalls) in global cellular networks. [DNS11] moves beyond coordinations among BSes within one operator network. It lets mobile users choose operators based on the application demand. [WJP13] investigate caching in backhaul cellular network to speed up content delivery.

2.4 State-of-the-art on Non-Cellular-Network

In addition to cellular network infrastructure, four modules are indispensable to enable mobile services, as shown in Figure 1.2. They are the front-end mobile, the back-end services, the intermediate wired Internet and the alternative WiFi wireless access. Many new research challenges also rise in these four areas as mobile devices and applications becomes pervasive. Here, we give a very brief summary of current work.

Mobile Mobile devices have been an extremely active research platform in the past decade. A large variety of problems have been addressed in the mobile context. They include localization [CBA10, LGY12], sensing [MRS09, RLL12], human-device interaction [YSC11], energy efficiency [QWG11], security and privacy [Lea05, CWY07, FFC11, Wal11, SZZ11, CFG11], offloading [BMV10, LLY10], optimization [QQH12, KLC12, Nat12] and new emerging mobile apps [CSB10, QBR11, ZCC12], to name a few. [CBA10] introduces social activities for human localization; [LGY12] integrates multiple localization techniques for phone localization; [MRS09] gives a recent report

of participatory sensing. [RLL12] introduces Medusa, a programming framework for crowd-sensing. [YSC11] exploits phone speakers to determine if the request comes to a car driver. [QWG11] profile energy use for mobile applications. Security study focus on various types of mobile malwares, including virus [Lea05, CWY07], SMS/call spams [FFC11], phishing [Wal11], Trojan [SZZ11], and intrusion on inter-application communication [CFG11]. [BMV10, LLY10] explore the feasibility and design to utilize WiFi to offload cellular network data traffic. [Nat12] supports continuous context-aware applications while mitigating sensing costs for inferring contexts. [QQH12] explores the opportunities and challenges for mobile caching. [KLC12] exploits network coding for video streaming on smartphones. [CSB10] describes SMS-based (140 bytes) mobile search; [QBR11] enables image tagging on smartphones; [ZCC12] designs a phone-to-phone action gaming.

Online services Most online services are moving toward cloud computing [VRC08] from the old client-server fashion. Cloud computing enables the use of computing resources and online services provided in a remote location (called as “cloud”). Such resources are usually offered by the underlying data centers via a virtualization technique. Recent years have witnessed a huge amount of studies in the cloud computing system and its underlying data centers. Here, we give a glimpse of several representative work.

In the clouds, many new designs are proposed to address classic system problems, since the cloud can be regarded as a super-large scale of distributed systems. For example, Google designed and implemented GFS [GGL03] as its file system, MapReduce [DG04] as its programming model, BigTable [CDG06] as its storage system, and Spanner [CDE12] as its distributed database. Afterwards, many following-up variants, such as Map-Reduce-Merge [YDH07], flat datacenter storage [NEF12] are proposed. Additionally, many research efforts also have been cast to address other technical concerns, such as privacy [GLS12, LWG13], security [HAR10, KCZ12], and performance [NKN12]. On the data center side, networking problems range from the

topology [ALV08, GHJ09, GLL09, SHP12, HKP11], routing [ACR10], TCP [AGM10, VHV12, PDH12] to the application layers [GRT10, LSY11, CZM11]. In particular, various networking topologies for data centers have been proposed by FatTree [ALV08], VL2 [GHJ09], BCube [GLL09], Jellyfish [SHP12], and wireless data centers [HKP11] via 60GHz flyways. [AGM10] addresses the problem of TCP-incast while [PDH12] exposes the TCP outcast problem and [VHV12] proposes a deadline-aware TCP. [GRT10, LSY11, CZM11] study the problem of bulk data delivery in data centers. Recently, soft-defined networking (SDN) primitives have been proposed and applied to data centers. It allows a programmable, centralized control of network traffic that decouples from the data plane [MAB08, RFR12].

Internet The Internet is a global system of interconnected computer networks to serve billions of users worldwide. It has had a revolutionary impact on our daily lives, since the Internet protocol suite (TCP/IP) standardization in the 1980s. The Internet has been the most successful networking system so far. However, driven by new user demands and inherent limitations in the Internet, the research community is moving forward to the next-generation design, such as 4D [GHM05], Named Data Network (NDN) [JST09] and XIA [HAD12], or addressing several challenging tasks, e.g., accountability [ABF08], content delivery and video streaming [TAY12].

WiFi WiFi is an alternative wireless access for mobile devices. It usually offers higher bandwidth than cellular networks but the connectivity is *local*. The re-association is required once it moves out of the coverage of the original access point. In industry, WiFi is evolving from the legacy 802.11a/b/g to 802.11n, which uses multiple antenna techniques and offer up to 600 Mbps bandwidth [Sta09]. In research, many advanced techniques are applied to boost wireless performance, such as, using white space (open spectrum) [BCM09] for higher bandwidth, exploiting fancy physical techniques available on soft-defined radio (SDR) [VBJ09, GK], developing duplex wireless transceiver [JCK11], application-specific optimization ([AK11] for video streaming).

Note that, these four modules are relatively independent of cellular network infrastructure; their challenges can be addressed separately. We believe that it can better support mobile service using an holistic approach(e.g., coordinating mobile and cellular networks, or making cloud service adaptive to heterogeneous mobile devices), but in this dissertation, we focus on the study of cellular network infrastructure only.

CHAPTER 3

Green Cellular Infrastructure

In this chapter, we address the problem to green cellular base station infrastructure, the most critical subsystem that consume a dominant fraction (about 80%) of overall energy [FZ08]. We first describe the problem of current 3G networks from the energy consumption perspective and identify two root causes for energy inefficiency. One is that each BS is not load-adaptive based on the BS power measurement; The other is that traffic load exhibits wide-range fluctuations both in time and over space. More trace analysis shows traffic and deployment diversity which inspires the GreenBSN design to green BS networks. We then present the design, theoretical analysis, implementation, evaluation of the GreenBSN solution.

3.1 Non Energy-Proportional BS Networks

We seek to make cellular BS infrastructure more energy efficient. Our goal is to design an *energy proportional* BS network. Energy proportionality (EP) seeks to achieve load-adaptive energy consumption in communication or computing subsystems [BH07, GS03]. The less traffic load the network carries, the smaller amount of energy the system consumes. The real problem is that, the current BS network operation is not energy




Figure 3.1: Energy-load curves for BS network in Region 1.

proportional to its carried traffic load.

Our study of real traces of 3G networks shows, the used energy is unproportionally large under light traffic load. This can be demonstrated by a case study on a regional 3G network in a big city with 177 BSes. Figure 3.1 plots the total consumed energy over a time window (here, 15 minute) ¹ versus the aggregate traffic load. The traffic load represents the total number of users who request for service in 15 minutes. The real power consumption is calculated based on power models described in Chapter 3.2. We observe that, even with light traffic (say, 2000 or below), the current energy consumption is still quite significant, about 95 KWh in total, approximately 95% of the peak power. In contrast, the desired (ideal) energy proportional operation should consume much less energy, about 5 KWh in total, under light traffic.

We have also digged into the trace and discovered why. It turns out that the traffic load at each BS varies significantly over time (see Figure 3.6(a) for a snapshot of traffic at four BSes in different regions). There is a large fraction of time (more than 10 hours over each 24-hour period) that the BS carries very light traffic. On the other

¹The energy is the power averaged over a time window (here., 15 minutes). We do not differentiate between power and energy hereafter.

hand, each BS system is not energy proportional to the traffic load. The root cause is that the large fraction (more than 50%) of energy spent on cooling power and idle-mode signaling and processing, are invariant of traffic load (as we show in Chapter 3.2), further contributing to non-energy-proportionality feature at each BS. Therefore, without energy-proportional operations, the 3G network suffers from large energy waste.

3.1.1 Design Goals and Challenges

Given that 3G network is not energy proportional to traffic load, our ultimate goal is to build a load-adaptive solution to energy savings in operational 3G networks. We design *GreenBSN*, a solution that approximates an energy-proportional 3G system using non-EP BS components, in order to cope with temporal-spatial traffic dynamics. Specifically, we aim to achieve three concrete goals while addressing their corresponding challenges:

1. The first goal is to achieve system-wide EP in the BS network. The corresponding challenge is how to approximate an EP system using non-EP BS components.
2. The second goal is to ensure negligible performance degradation while achieving energy efficiency. The corresponding challenge is how to meet location-dependent coverage and capacity requirements.
3. The third goal is to build an EP solution compliant with the standard specification. The corresponding challenge is how to leverage the existing 3G mechanisms for the purpose of energy efficiency.

To achieve these goals, we completely power off under-utilized BSes when their traffic load is light and power them on when the load becomes heavy. The challenge is how to devise a *distributed, practical* solution that uses a small set of active BSes, while satisfying three requirements of traffic capacity, communication coverage, and minimal on/off switching of each BS.

3.1.2 Roadmap to the Solution

To build the GreenBSN solution, we first need to quantify the root cause to non-EP operations. That is,

Q1. What are the BS power consumption model and critical factors?

Q2. What are the characteristics of traffic load in operational 3G networks?

We use real power measurement, traffic traces and BS deployment map to conduct detailed analysis on power model and traffic dynamics (Chapters 3.2 and 3.3). The study of diversity both in time and over space sheds light on the GreenBSN design. We next propose the GreenBSN design and implement. In particular, we address another two issues:

Q3. Given the traffic dynamics, how can we achieve network-wide energy proportionality (EP) using non-EP BS components?

Q4. How can the proposed solution work with the current 3G standard?

We take a location-dependent profile-based approach. The three key components are grid-based BS clustering, traffic envelop profiling and graceful on/off selection. In particular, we divide the network into grids, so that BSes in each local grid can replace each other when serving user clients. We then perform location-dependent profiling to estimate the aggregate traffic among BSes in the grid. Based on the peak/idle of the traffic profile, we decide the corresponding set of active BSes for each duration. It turns out that, if we select the active sets appropriately, we only need to power on a sleep BS and shut down an active BS at most only once during each 24-hour period. To make GreenBSN standard compliant, we reuse cell breathing, handoff and RNC control techniques which have been enabled at existing BSes.

3.2 Understanding BS Power Consumption

We investigate and model the energy characteristics of base stations in cellular networks to obtain realistic parameters to understand the key factors to energy efficiency and evaluate energy saving in our scheme.

Figure 3.2 shows a typical BS in 3G UMTS networks. Each BS is the physical unit providing network access services to mobile users via its air interface. It has the communication subsystem and the supporting subsystem. The communication subsystem includes Remote Radio Unit (RRU), Base Band Unit (BBU), feeder and RNC transmission component. RRU is the radio-specific hardware for each sector. BBU provides all other communication functions, including control, base band, switching and Iub interfaces to RNC. Each BS may install several RRUs and BBUs. Feeder is the optical-fiber pair cable that connects RRUs to BBUs. The supporting subsystem includes the cooling subsystem, the power supply module and other auxiliary devices (e.g., for lighting and environment monitoring). The cooling subsystem, including air conditioning and fans, retains an appropriate operation temperature at the BS.

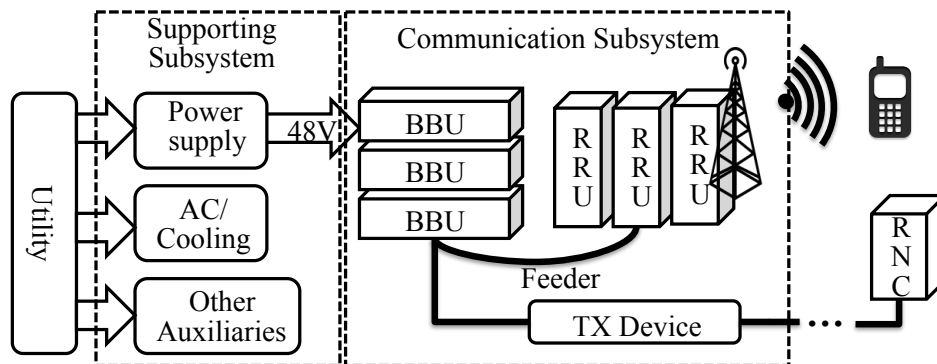


Figure 3.2: A typical base station in 3G networks.

We now model the overall BS power consumption, including both the radio communication and the supporting parts. We use real measurement data taken on both transmission and cooling systems at BSes. The cooling is the dominant power consumer

in the auxiliary subsystem and thus this measurement approximate (though underestimate) the overhead power for the supporting subsystem. The total power consumption P at a BS is given by $P = P_{tx} + P_{misc}$, where the first part P_{tx} accounts for power used for the communication subsystem, and the second part P_{misc} records the auxiliary power for cooling, power supply and monitoring. We next show that P_{tx} mainly changes with carried load while P_{misc} typically remains constant given a fixed operating environment.

Modeling P_{tx} Using real measurement data on transmission power, we find out that linear models can offer reasonably good approximation for a variety of BSes; This model has also been widely adopted in the literature [ARF10,RFF09,DBM10,BNC10]. Figure 3.3 gives the scatter plot of measured power at three BSes, with regards to their carried load. The figure clearly shows that a linear model can approximate the transmission power with respect to the carried traffic load, i.e., $P_{tx}(L) = P_{\alpha} \cdot L + P_{\beta}$, where L is the traffic load².

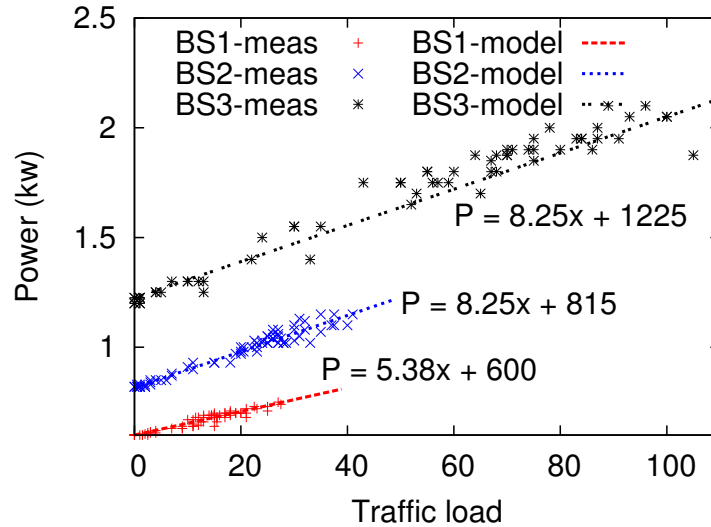


Figure 3.3: Measurements and model estimates of BS transmission power (P_{tx}) with regards to their carried loads.


²Here, the traffic load is measured by the number of user requests in a 15-minute bin. We also examine the traffic load in terms of the number of occupied traffic channels (TCHs) or utilization level (i.e. a traffic load factor) later. It still approximates a linear model but the value of model coefficients change.

The above empirical model can be also explained by the actual BS operations. The two dominant components in P_{tx} are the power consumed by RRUs and BBUs. When the traffic load is heavy, RRU has to spend more power to support more active links. Therefore, it increases in proportion to the traffic volume. On the other hand, BBU does baseband processing for all frequency carriers used by the BS. No matter how many links are active, its power consumption is mainly determined by the number of frequency carriers unless it is in sleep mode.

Note that the power coefficients (i.e., slope and offset) may vary over BSes. This is caused by different products and the changing number of installed BBUs and RRUs at each BS. The measurement and model also match with BS product data sheets, for example, three popular BSes have P_{tx} ranging from 600w to 3000w [Eri, Hua, Mot]. In our model, transmission power also increases when the operational range expands. Specifically, when the BS reaches its maximum transmission range via cell breathing or duplicate long-range radios (see Chapter 3.4 for details), we model that the power also grows in proportion to the traffic load but uses a larger coefficient P_a , say, P_a doubles at its maximum range. Our design and evaluation consider such diversity factors.

Modeling P_{misc} We focus on modeling the cooling (i.e., air conditioner) power consumption since it is the dominant factor in P_{misc} based on real measurement. Previous work does not model this part, though it is known that cooling may consume about 50% power at BSes [FZ08]. The cooling power consumption depends on the amount of the extracted heat and the desired operating temperature. It also varies with chillers that use a variety of compressors and drivers.

Figure 3.4(a) shows the scatter plot of the cooling power and temperature at a BS in 2010. It is seen that the cooling power mainly depends on the temperature. It increases approximately linearly from 1000w to 2000w when the environment temperature varies from $-10^{\circ}C$ to $30^{\circ}C$ (i.e., from winter to summer). We also check daily and yearly pattern in Figure 3.4(b). The upper line is the yearly pattern that varies with four seasons. The lower line is a 4-day measurement in early winter. It shows though the cooling



(a) Power vs. temperature

(b) Power vs. time, two scales of 4-day and 4-season

Figure 3.4: Measurement of the cooling power (P_{misc}) at one BS in 2010.

power fluctuates slightly at different hours of a day (e.g., BBUs and RRUs tend to raise the air temperature and the chiller workload), it can still be approximated as a constant within a short period of time (say, a day), here in [1200w, 1400w]. Over a larger time window (say, a year), it varies with the external environment temperature. For simplicity, we assume p_{misc} remains constant on a daily basis but changes with seasons.

In summary, the power model can be represented by

$$P = P_{\alpha} \cdot L + P_{\beta} + P_{misc} = \omega \cdot L + \theta. \quad (3.1)$$

We model the total BS power as a linear function of its expected load, with ω being power efficiency constant (used power for the unit load), and θ and the cooling power as a linear function of its expected thermal load, depending on external air temperature and base station load. Specifically, at a given BS, in short term (say, several days), P_{misc} can be approximated as a constant (use the median in a day, 1300w in our example) and in the long term, it may fluctuate in different seasons. Table 3.1 presents six power models to be assessed. The first five models are homogeneous and the last one assesses the tradeoff between high capacity and high energy efficiency, where BSes with larger capacity consume more power.

Our model is consistent with other power models used by the literature [HGX11, DBM10, ARF10]. Moreover, this model extends early efforts and develop a single unified model. Previously proposed models only consider radio transmission but ig-

#	Model	Setting	Description
1	Sp/Fall	$P = 6L + 1500 + 600$	Normal P_{tx} with medium P_{misc} , in general cases (benchmark)
2	Winter	$P = 6L + 1000 + 600$	Normal P_{tx} with low P_{misc} , usually in cold seasons or using high-efficient cooling technique, e.g. fanless cooler or cooling using natural resources
3	Summer	$P = 6L + 2000 + 600$	Normal P_{tx} with high P_{misc} , usually in hot seasons or using old-fashion cooling technique
4	Sp/F-low	$P = 4L + 1500 + 600$	Low P_{tx} with medium P_{misc} , using novel hardware designs, e.g. energy efficient power amplifier, or deployed in urban or indoor cases with small capacity and coverage
5	Sp/F-high	$P = 8L + 1500 + 600$	High P_{tx} with medium P_{misc} , using previous generation hardware, or deployed in rural regions with high capacity and coverage.
6	Hybrid	H, $P = 8L + 1500 + 800$ M, $P = 6L + 1500 + 600$ L, $P = 4L + 1000 + 400$	P_{misc} increases as P_{tx} , with more cooling load in case of larger P_{tx} . Consider the base station configuration diversity.

Table 3.1: Different power models used in this work.

nore power for cooling and other auxiliary devices [DBM10, BNC10], or over/under-estimated power consumption coefficients [RFF09, ARF10]. It shows that the cooling subsystem and transmission module consume a significant portion (more than 50%) of overall power at each BS, regardless of the traffic load. This is the dominant factor that leads to energy inefficiency for the 3G infrastructure. Recent component-based proposals [Eri08, EL04, Hua] can reduce the power overhead (θ) to some extent, but cannot eliminate it. That is to say, single BS is non-energy-proportional.

3.3 Characterizing Diversity in Trace Analysis

We now present our measurement results on 3G traffic dynamics and network deployment diversity. Our trace analysis shows that BS networks convey strong temporal-spatial diversity and under-utilized BSes are very common in practice. Thus, it results in huge energy waste when they are always-on in spite of traffic loads. In addition to

diversity characterization, we dig out design insights on how to improve the current 3G network’s non-energy-proportionality.

We use real traces collected from two large, independent, operational 3G networks in the world. The first operator dataset contains BS locations and 15min-bin traffic volume records in four regions for two months (August–October, 2010). The second operator dataset provides extra information on BS capacity and dynamic configuration, as well as one-week (Dec 15–21, 2011) hourly traffic and BS deployment in a much bigger region. We hide the detailed deployment map for privacy concerns and plot BS locations in these five regions in Figure 3.5; They have different geographic scales and represent diverse city types: Region 1 is a large, populous city, Region 2 is a medium-sized city, Regions 3 and 4 are large cities in a large metropolitan area, and Region 5 covers a hybrid geographic landscape, including a medium-sized coastal city, its wide suburb, and an island around it. All regions have diverse residential and downtown areas. The coverage area and the number of BSes in each region are given in Table 3.2.




Figure 3.5: Maps of base station locations in five regions. The right plot in Region 5 plots a 10km x 10km subregion (called as “5A”) where 281 BSes are deployed. Dotted rectangles A and B indicate residential and downtown areas.

	Region 1	Region 2	Region 3	Region 4	Region 5
Operator	OP-I	OP-I	OP-I	OP-I	OP-II
Area (KM×KM)	11×11	8×4	16×28	30×45	95×62
# BS	177	45	154	164	432
BS density	dense	dense/normal	normal/sparse	sparse	hybrid

Table 3.2: Basic statistics of 5-region traces.

We first studies traffic dynamics both in time and space using the four-region traces

from the first operator. We do not use the second dataset since the record time is too short; we use it as an independent source to verify the former findings on temporal-spatial traffic patterns. Moreover, we use it to explore capacity and utilization diversity, which are related to BS capacity settings that are not available in the first dataset.

3.3.1 Traffic Diversity Over Time

Based on our trace analysis, we make several key observations on the 3G traffic dynamics over time.

Finding 1: Temporal traffic dynamics We first find that each BS exhibits high traffic dynamics over time. Figure 3.6(a) plots the traffic load at four individual BSes in different regions. We observe strong diurnal patterns on both daily and weekly basis, alternating between peak and idle durations³. We separate the weekday and weekend data here, and only present the weekday case unless explicitly stated; the result for weekend is similar.

To quantify the degree of temporal traffic dynamics, we compute the ratio of peak-to-idle traffic load at each BS in four regions. We define the peak (/idle) duration of each BS as the hour h , when it has the maximum (/minimum) traffic load (typically between 10AM-18PM for peak, or 1AM-5AM for idle), plus two adjacent hours, i.e., $h - 1$ and $h + 1$. Figure 3.6(b) presents CCDF of peak-to-idle traffic-load ratios in four regions. We see that the peak-to-idle traffic ratio is larger than 4 in most (70-90%) BSes, and the smaller ratio (say, 2-4 in Region 1) is only due to relative small traffic volume at BSes. We also study the effect of time window size (here, 3hr) and find that large peak-to-idle ratios still exist when the window is smaller than 8 hours.

Insight 1. *This result shows that the traffic distribution of each BS is quite diverse over time everywhere. Such strong temporal diversity indicates the under-utilization of each BS in the time domain, resulting in system-wide energy inefficiency.*

³We use the term “idle” duration for light traffic cases in our work.

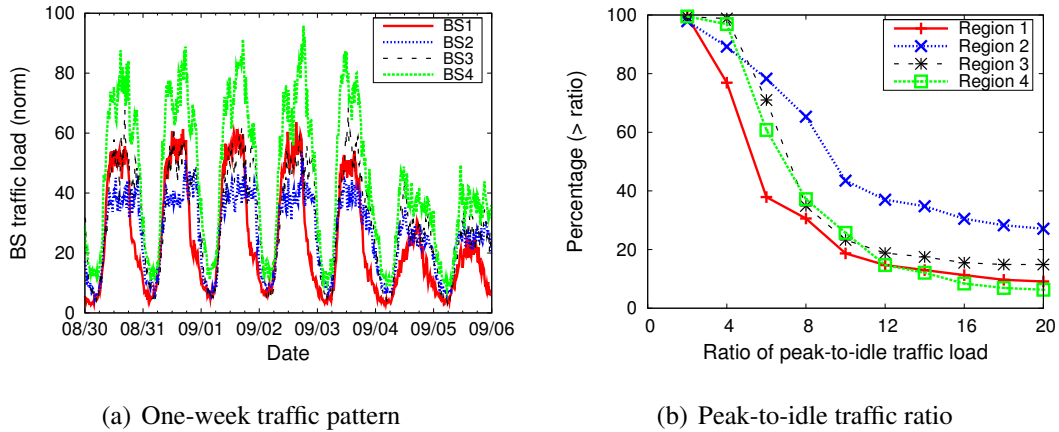


Figure 3.6: Traffic dynamics over time.

Finding 2: Near-term traffic stability We also observe that the traffic volume be stable over short term (e.g., the same time of consecutive days), while it may slowly evolve over a long term (e.g., 26% global increase in 2010 [Cis11]). Although the traffic load fluctuates over time, the time-of-the-day traffic at each BS is quite stable over consecutive days (see Figure 3.6(a)). For example, BS1 has similar traffic load at 5 pm in Days 1 and 2, Days 2 and 3, and so on.

Our results show that, in all four regions, the autocorrelation values are higher than 0.963 for 70% BSes – confirming strong correlation between traffic load during two consecutive days. To measure such near-term stability, we compute the near-term traffic variation $V(i, t)$ at time t at BS i :

$$V(i, t) = |R(i, t_{cur}) - R(i, t_{prev})| / R(i, t_{prev}), \quad (3.2)$$

where $R(i, t_{cur})$ and $R(i, t_{prev})$ denote the traffic load of BS i at time t on the current day and on the previous day. Table 3.4 shows the near-term variation statistics using our two-month data in four regions. We see that, at any time, the traffic load difference in two consecutive days is less than 20% for 70% BSes. We also note that high variation values are mostly caused by the low traffic volume at the idle time and their absolute values of traffic difference are, in fact, quite small. We further examine the impact of different aggregation granularity (e.g., from 15-min bins to several hours). The near-

term variation increases as the aggregation granularity grows, but remains highly stable when the window is smaller than 2 hours.

Traffic load auto-correlation at one-day lag						Traffic load variation in consecutive days					
Location	10th	30th	50th	70th	90th	Location	10th	30th	50th	70th	90th
Region 1	0.831	0.919	0.950	0.963	0.979	Region 1	2.1%	6.7%	12.1%	20.1%	38.7%
Region 2	0.833	0.932	0.957	0.972	0.980	Region 2	1.9%	6.4%	11.8%	20.6%	45.0%
Region 3	0.898	0.938	0.954	0.966	0.977	Region 3	2.1%	6.8%	12.3%	20.4%	42.0%
Region 4	0.907	0.940	0.959	0.969	0.983	Region 4	2.0%	6.3%	11.4%	18.8%	36.6%

Table 3.3: Traffic autocorrelation with 24-hour lag. Table 3.4: Traffic variation in consecutive days.

Insight 2. *The near-term stability result makes a case for traffic profiling to estimate/predict the next day’s traffic trend and motivates us to develop power-saving schemes based on traffic profiles. The measurement also indicates that the hourly traffic aggregation achieves good balance between estimation accuracy and simplicity.*

Finding 3: Time-domain multiplexing diversity We find that the aggregate traffic load in a region hardly reaches the aggregate BS capacity in the region. To verify such a trend, we define time-domain “multiplexing” gain $M(t)$ as the ratio of the sum of the peak traffic at each BS (i.e., lower bound of BS capacity) to the aggregate traffic load at time t in the region: $M(t) = \sum_i R(i, t_{max}) / \sum_i R(i, t)$, where $R(i, t)$ is the traffic load of BS i at time t ; t_{max} is the peak traffic time.

Figure 3.7(a) plots the multiplexing gain $M(t)$ in four regions. We see that the multiplexing gain is 4-15 at midnight and around 2 (ideally, it should be close to 1) even during daytime in all regions. Note that the gain can be even larger in reality because the operators often deploy BSes with much larger capacity than the actual traffic demand to account for forthcoming market growth. The root cause is that not all BSes reach their peak load simultaneously. We study the peak hour distribution in subregions A (residence area) and B (business area). For example, the peak hour spans from 10 AM to 6 PM in subregions A (residence area), and from 4 PM to 8

PM in B (business area), as shown in Figure 3.7(b). The operator has to deploy the infrastructure that can accommodate the peak traffic at each location, even though the peak load may only last two or three hours a day. As the peak hour varies with each location, the deployed capacity (i.e., the sum of each BS’s capacity) is much larger than the actual traffic volume at the time. Our following study shows that the operator may over provision BS capacity so that the multiplexing gain using aggregate capacity is even larger.




Figure 3.7: Illustration of temporal multiplexing diversity.

Insight 3. *This multiplexing gain shows that the aggregate BS capacity is highly under-utilized in each region. The inherent temporal-spatial diversity opens venue for energy savings via aggregating traffic load at BSes.*

3.3.2 Traffic and Deployment Diversity in Space

Finding 4: Diverse BS deployment density The BS deployment density varies across locations (see Figure 3.5 for location distributions). In the hot spots of a city (e.g., subregion B), more BSes are provisioned, thus creating location-dependent diversity. Figure 3.8(a) depicts the distribution of the number of neighbors per BS (within 1 Km), representing the BS deployment density in four regions. We see that the deployment density is quite diverse across different regions, as well as in the same region. For instance, Regions 1 and 2 have more dense BS deployment than Regions 3 and 4. We

also see that a large number of BSes have multiple neighbors, especially in Regions 1 and 2. For example, for more than half of BSes in Region 1, each has at least 10 neighbors within its 1Km range. In contrast, Region 4 has the most sparse deployment; only 40% BSes have multiple neighbors.

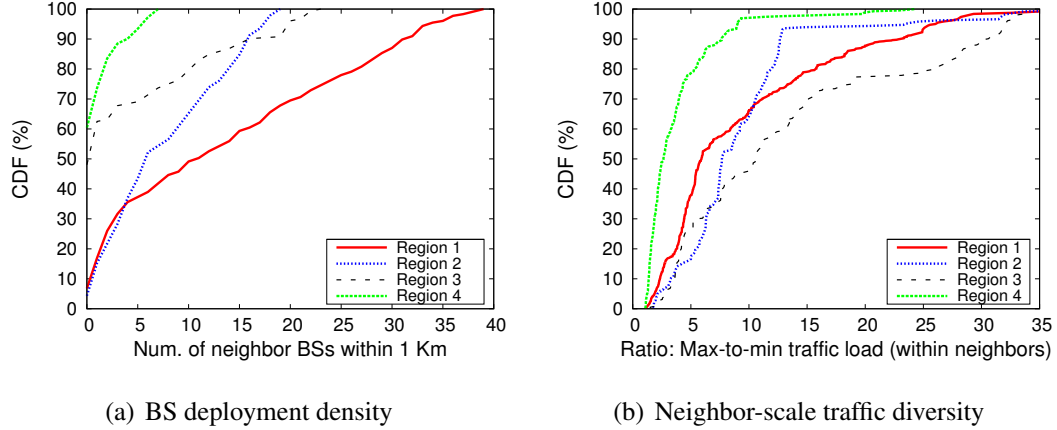


Figure 3.8: Spatial diversity in deployment and traffic.

Insight 4. *This BS deployment practice provides us an opportunity to exploit such topological “redundancy” for energy-proportional power savings, and the expected gain tends to vary across regions.*

Location	Peak time			Idle time		
	20th	50th	80th	20th	50th	80th
Region 1	4.0	6.5	17.1	2.3	3.9	6.0
Region 2	6.2	9.1	12.3	3.7	6.2	8.7
Region 3	7.5	13.4	28.2	3.6	4.9	10.6
Region 4	1.7	3.2	6.0	1.5	1.9	3.1

Table 3.5: Max-to-min traffic ratio in neighborhood

Finding 5: Spatial traffic diversity Another key observation is that traffic load intensity is quite diverse even in each local neighborhood (i.e., traffic loads among the closely located BSes). Figure 3.8(b) shows the spatial traffic diversity among neighboring BSes. Each point represents, at any given time of the day, the traffic-volume ratio

of the maximum-traffic BS and the minimum-traffic BS within 1 Km range of each BS in four regions. We see that the max-to-min traffic ratio is larger than 5 in 50% cases, and larger than 10 in 30% cases (in Regions 1, 2, and 3)⁴. We also observe that such neighborhood-scale spatial traffic diversity is more evident during the peak time. Table 3.5 presents the max-to-min BS traffic ratio at peak and idle times. Note that, for example, the spatial diversity at the peak time becomes a factor of 13.4 in 50% cases in Region 3.

Insight 5. *Such strong neighborhood-scale traffic diversity indicates the under-utilization of a group of BSes in the spatial domain, and sheds lights on energy savings in each local area.*

3.3.3 Capacity and Utilization Diversity

We now use the second dataset to study deployment diversity, particularly regarding BS capacity. We first assess traffic dynamics in Region 5, and confirm that the above findings still hold for this independent dataset. For example, Figure 3.9 plots six-day traffic loads at four nearby BSes. It shows that they still convey strong diurnal patterns (Finding 1). The slight difference from the previous study is that there is a clear slump around noon and peak traffic often occurs at night (6PM-10PM); Traffic patterns are determined by local human activities; in fact, this city is well known its dining and entertainment activities at night. This example also demonstrates its spatial traffic diversity; traffic load fluctuates at neighboring BSes. We cannot validate its near-term traffic stability (Finding 2) since we only have one-week data. However, we do observe high multiplexing gain (Finding 3) and diverse BS deployment (Finding 4).

Finding 6: Spatial capacity diversity The BS capacity is measured as the number of traffic channels (TCHs) that are logical channels to accommodate voice or data traffic [3GP12b]. In the 2G GSM or 3G UMTS context, the number of TCHs is mainly

⁴Region 4 gives the lowest max-to-min traffic ratio due to its sparse BS deployment.

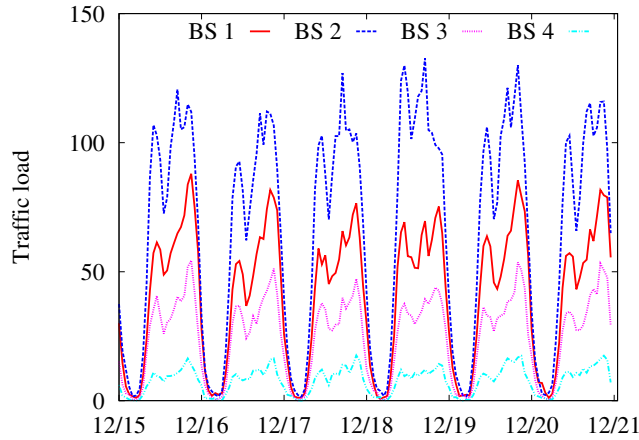



Figure 3.9: An example of six-day traffic traces of four BSes in the second data set.

determined by the number of frequency carriers assigned to a BS. It depends on the physical hardware setting, such as how many frequency carriers RRUs can handle for RF signal transmission/reception and how fast BBUs can process baseband signals. For example, one macro BS can support up to twelve carriers while a micro BS may use at most two carriers. To easily quantify BS utilization, the traffic load in Region 5 uses the number of occupied TCHs at one time. The temporal-spatial dynamics is consistent with the former results using the number of users requests as the traffic load.

Figure 3.10(a) plots the map of BS capacity in subregion 5A. Each Voronoi diagram represents the coverage of one BS [Aur91]. It clearly exemplifies that BS capacity is highly diverse over space. Figure 3.10(b) plots the probability distribution function (PDF) of BS capacity in Region 5 and Subregion 5A. It shows that BS capacity is almost uniformly distributed in a small range (here, (50, 200) in Region 5) and the probability greatly reduces beyond that. That is, BSes with super high capacity are rare. In the densely-deployed area (e.g., Subregion 5A), capacity is comparably higher; For example, the most common BS capacity setting appears in (200, 250), much larger than the average in the entire region. This matches well with our expectation since the traffic load is usually heavier in dense-deployment areas. We also notice that BS capacity configuration is discrete in Subregion 5A; It implies that it might not to be

feasible to configure BS capacity continuously or arbitrarily. BS capacity options are also constrained by BS device models.

Insight 6. *Heterogeneous capacity further signifies the opportunity to utilize more capable BSes for energy savings. Fewer BSes with higher capacity can afford equivalent aggregate capacity by a larger number of less-capable BSes.*



(a) Capacity in Subregion 5A

(b) Capacity distribution

Figure 3.10: Illustration of BS capacity diversity in Region 5.

Finding 7: Temporal-spatial utilization diversity We next characterize BS utilization in time and over space. The BS utilization is defined as the ratio of its traffic load to its capacity. Given that BS capacity is almost constant (i.e., configuration is done during radio planing), each BS is designated to accommodate its peak traffic load, except load surge caused by rare, special events (e.g., celebration parade).

Figure 3.11(a) shows the probability density function (PDF) of peak utilization; it is normalized by setting the maximal peak utilization as one. We observe that, most BSes are not fully utilized; the peak utilization at most BSes is only within the range of (0.4, 0.8). Such utilization levels do not happen without rationale. One is to reserve resource so as to maintain a low service blocking rate. Traffic requests (e.g., voice calls) are blocked if there is no TCHs available at that time. Based on various traffic models such as Erlang models [Wik, HT07], cellular operators reserve enough capacity room




Figure 3.11: Peak and temporal utilization in Region 5 and 5A.

under normal usage. For example, given 1% block rate, the BS with 100 TCHs can afford 84 Erlang load (one Erlang load is defined as the accumulative traffic that fully occupied one TCH for a unit time). Second, extra capacity is probably reserved for rare surge events and the upcoming usage growth.

To understand temporal utilization, we compute the hourly utilization at each BS (hourly traffic over its capacity) and plot its PDF in Figure 3.11(b). It takes both temporal and individual factors into account, and demonstrates the overall utilization in this region. It shows that, BS networks experience low utilization under temporal dynamics (also seen in Finding 1). In most cases, the utilization is smaller than 50%, and the most common utilization is as low as 5%.

Insight 7. *Utilization dynamics, especially low utilization, implies that there exists enough room to further improve system-wide energy efficiency.*

3.4 GreenBSN: Towards Green BS Networks

GreenBSN is inspired by all seven findings. If traffic varies over time (Insight 1) and in space (Insight 5), consuming energy adaptive to the traffic variation becomes critical. The near-term traffic stability (Insight 2) makes a case for *profile*-based approach to estimating traffic envelope at any time. Since the multiplexing gain is high (Insight 3),

the profile-based scheme on aggregate traffic, rather than individual BS traffic, will deem more effective. To leverage deployment diversity (Insight 4), we power off under-utilized BSes when their traffic is light and turn them on under heavy traffic. Capacity diversity (Insight 6) allows to power on fewer capable BSes to further reduce energy consumption. The extra capacity that is not fully utilized during most times (Insight 7) implies the applicability of the power-off scheme.

The overall design of GreenBSN takes a grid-based, location-dependent profiling approach. We divide the network into grids, so that BSes in each grid cell can replace each other when serving user clients. Once the grid is established, we perform location-dependent profiling, which estimates the traffic envelope for the aggregate BS traffic in the grid. Given the peak and idle hours of the traffic profile, we decide the corresponding set of active BSes for each duration. It turns out that, if we select the sets appropriately, we only need to power on a sleep BS and shut down an active BS only once during each 24-hour period. This minimal on/off switching works well with the cooling subsystem, which needs 10s of minutes when adjusting to the desired operating temperature inside each BS upon power-on. We inject diversity at certain BSes to provide higher room to adjust its coverage in order to increase the chance to power off neighbor BSes. This way, GreenBSN offers a distributed solution that uses a small set of active BSes. It also satisfies three requirements of traffic capacity (i.e., traffic does not exceed BS capacity), communication coverage (i.e., each location is covered by at least an active BS), and minimal on/off switching of each BS (i.e., we avoid powering on/off each BS frequently). We next elaborate our design.

Our design also eliminates several limitations of the popular optimization-driven approach [DBM10, MCC09, BNC10]. These drawbacks include a centralized rather than distributed scheme, approximation to the optimal solution, excessive on/off switching, unrealistic BS power consumption model, one-time optimization targeting instantaneous traffic load, and difficulty in addressing deployment diversity and node heterogeneity. Finally, they do exploit multiplexing gain to minimize active BSes.

3.4.1 Grid-based BS clustering

The grid-based profiling approach estimates traffic in a given area. To this end, we devise a grid-based profiling scheme that partitions the 3G network into virtual grids and profile aggregate traffic within each grid. To make this scheme work, we It addresses two issues: (1) How to determine the grid to partition the network and facilitate powering off under-utilized BSes in a given area? The proposed solution has to accommodate diversity in node deployment and communications. (2) How to perform location-dependent traffic profiling to exploit the multiplexing gain over time and among local BSes? We now elaborate our solution to these two issues.

The grid-based BS clustering partitions the BS network into virtual grids. We can then power off under-utilized BSes in each grid. The design is motivated by the example of six BSes in Figure 3.12, where r_i and R_i denote the normal and maximal communication ranges for BS i , respectively. No matter how large the capacity active BSes 5 and 6 have, the left area loses access coverage when three BSes 1, 2, and 3 power off. Therefore, sufficient capacity is needed at each location. We devise a grid-based approach to decoupling coverage requirement from capacity demand.

We partition virtual grids so that BSes in each grid are equivalent. BSes are *equivalent* if they can replace each other when communicating with user clients. We use location information and transmission range of each BS to decide whether BSes in spatial proximity are equivalent or not. Location coordinates can be obtained by GPS or other systems when operators deploy their infrastructure. Transmission range of a BS may vary from 200m to 1km in cities and from 1km to 5km in rural areas [Mis04]. It can be different among BSes due to antenna configuration and placement, transmit power and environment. For example, the antenna is often placed at the top of a high mast (say, about 40-80m) to obtain a large 5km transmission range at a rural site.

Two BSes are equivalent if they can cover each other in data communication with clients. Specifically, let $d(i, j)$ be the distance between two BSes i and j , then they are

equivalent if

$$r_i + d(i, j) \leq R_j, \quad r_j + d(i, j) \leq R_i. \quad (3.3)$$

Note that the above procedure can accommodate diversity in node deployment and communications, e.g., the changing distance $d(i, j)$ between any pair of nodes i, j . In the example of Figure 3.12, BS 1 is equivalent to BSes 2 and 3, but is not equivalent to BS 4.




Figure 3.12: Example of virtual grids. Left: BS on/off status. Right: virtual grids.

A virtual grid cell is formed when all the BSes in it are equivalent. Once a BS is not equivalent to every BS in the current grid, we create a new grid cell. Since grid formation can be nonunique, we use a simple heuristic “northwest rule” to decide our grid construction. We start from the northwest corner in the BS deployment map (i.e., top-left corner), cluster all equivalent BSes from top to down and from left to right, and generate a new grid-cell when a BS is found to not be equivalent to at least one BS in the current cell. We repeat the process until we exhaust all the BSes in the 3G network. In the illustrative example of Figure 3.12, three grid cells are thus formed following this rule. We notice that formation along other directions may generate different virtual grids, but would not much affect energy savings. No matter what grids are created, it does not change the inherent proximity. Close nodes belong to the same grid with high probability. For example, if the grid is formed in “northeast” rule (i.e., top-right first), we have another three grids: 6 and 5, 4 and 3, 2 and 1. Each virtual grid still has similar redundancy (the average density is 2 here) and offers local capacity at slightly different spots. We note that the grid-based clustering does not allow across-grid cooperation for

energy savings and incurs sub-optimality. We discuss it in Chapter 3.5.

3.4.2 Location-dependent Traffic Profiling

We now devise a profiling scheme that estimates the envelope of aggregate traffic demand in a local grid. We have two specific goals: (1) The solution should not underestimate the traffic so that the miss rate to serve clients is undesirably high. In the meantime, it should not be too conservative to overestimate the traffic by an unattainable margin. (2) It should be able to leverage the multiplexing gain of local traffic over time and among BSes in the grid.

To exploit traffic multiplexing over time and among local BSes, we profile the *aggregate* traffic in a grid cell. Finding 2 states that traffic is highly stable and predictable between consecutive days during weekdays and weekends. In consequence, we divide each day into 24 hourly intervals, compute the statistics of each hourly interval, and derive the traffic envelope for the given hour (in fact, all <2-hour intervals have similar and good performance). We differentiate a weekday from a weekend day, but treat all weekdays or weekend-days similarly. Specifically, for the i -th hour of k -th day that we stack together consecutive weeks, we compute the moving average $\bar{S}(i, k)$ and standard deviation $\bar{D}(i, k)$ as follows:

$$\bar{S}(i, k) = (1 - \alpha) \cdot \bar{S}(i, k - 1) + \alpha \cdot S(i, k), \quad (3.4)$$

$$\bar{D}(i, k) = (1 - \beta) \cdot \bar{D}(i, k - 1) + \beta \cdot |S(i, k) - \bar{S}(i, k)|, \quad (3.5)$$

$$EV(i, k) = \bar{S}(i, k) + \gamma \cdot \bar{D}(i, k), \quad (3.6)$$

where $S(i, k)$ is the hourly sample value of the aggregate traffic in the grid for i -th hour during the k -th day, and $EV(i, k)$ is estimated hourly traffic envelop. α, β are the smoothing parameters, chosen as $\alpha = \frac{1}{8}$ and $\beta = \frac{1}{4}$ in our prototype, and γ is a design parameter that offers a tuning knob to balance between tight estimate and miss ratio. Note that our goal is not purely to get the best estimation of traffic load, it is to provide traffic predication to make on/offdecision. The operators would rather be

conservative to turn on some unnecessary BSes than fail to serve mobile users because they underestimate traffic loads and turn off nodes that must be on. Our profiles is to find a cap that can cover the real traffic and it is also simple and accurate.

Note that our group-based profiling improves energy efficiency when traffic load is heavy, compared with the individual profiling. The individual approach estimates each individual traffic envelope first without extracting the multiplexing effect among local BSes. In contrast, our group-based profiling can improve energy efficiency when traffic load is heavy. Figure 3.13 shows an example of 15 BSes in one grid with several micro grids. The peak hours in two micro grids (marked by “+” and “x”) vary slightly and exhibit different patterns even within the same grid. As a result, it leads to about 5-8% energy-saving gain at peak hours when using the group profiling scheme (see Chapter 3.7.2 for details).

In case of sudden events, traffic can not be predicated by historical data. We can combine profiling with real time traffic monitoring and trigger schemes adaptive to real traffic load (discussed in Chapter 3.6). We also take the option to reserve a fraction (say, 10%) of the capacity in a BS when we make on/off decisions in worse case of unpredicted traffic transition.




Figure 3.13: Profiling examples.




Figure 3.14: Example of BS graceful selection.

3.4.3 Graceful Selection of Active BSes

Given the traffic profile in each grid, we next select the right set of active BSes and power off under-utilized BSes. The design has to reach two goals of minimizing the number of on/off operations and satisfying both coverage and capacity requirements. This requirement is critical to tuning to the right operating temperature by cooling subsystems before BBUs start to function at each BS. The air conditioning may take tens of minutes to adjust the operating temperature inside the BS to the desired degree by the BBU and other communication components. Consequently, it may take half an hour or more to fully activate a BS. To this end, our solution has three components: (1) selection of active BSes for the peak hour(s), (2) selection of active BSes for the idle hour(s), and (3) smooth transition between the idle and the peak.

Selection of active BSes for peak hour Given the 24-hour traffic profile at a given grid, we first find the hour(s) with heaviest traffic. For this peak duration, we need to select the set of active BSes in the grid, denoted by S_{max} . Based on the fact that the residual energy contributes a large percentage, we reduce the number of active BSes to save energy. On the opposite side, the *local*, aggregate capacity of all active BSes has to be large enough to accommodate *local* traffic. Our algorithm thus prefers the BSes with larger capacity. We rank all the BSes in the grid in decreasing order of their capacity values $C(BS_i)$, say, $C(BS_1) \geq C(BS_2), \dots, \geq C(BS_n)$. Then we select the smallest number m of active BSes so that $\sum_{k=1}^m C(BS_k) \geq EV_{max}$, where EV_{max} is the estimated traffic envelop at the peak hour. Then, the set of active BSes for the peak hour S_{max} is given by $S_{max} = \{BS_1, \dots, BS_m\}$. This heuristic ensures the minimum number of active BSes in the grid. When BSes have heterogeneous capacity (Insight 6), the smaller number of BSes with high capacity can afford the traffic demand. Assume that all local BSes use same power models, we can easily prove that the algorithm is optimal to ensuring minimum total energy in the grid (Lemma 3.1). When BSes have heterogeneous power models (i.e., different ω, θ), we will select the high-energy-

efficiency BSes with high priority if their capacity exceeds the traffic demand. We repeat the above algorithmic procedure for each grid, thus obtaining the active BSes for each grid during its peak hour.

We repeat the above algorithmic procedure for each grid in the network, thus obtaining the active BSes for each grid during its peak hour. Note that the peak hours in different grids may be different. Once active BSes are selected for each grid, our algorithm can meet requirements for both coverage and capacity. Note that two nodes in adjacent grids may cover each other. This offers new opportunity to further save power by merging active BSes in adjacent grids. However, our study shows that this option would add much higher complexity to the design; we trade marginal power savings for design simplicity in this work.

Selection of active BSes for idle hour When deciding the set of active BSes for the idle hour that has the smallest amount of hourly traffic in the 24-hour profile, we devise a different scheme. Rather than select active BSes from all candidates in the grid, we select the active set only from the superset S_{max} , which we have calculated for the peak hour. We can use similar selection policy to find the BS set for the idle hour, denoted by S_{min} . It is guaranteed to be a subset of S_{max} .

We use the above scheme to minimize the on/off switching of BSes by sticking to the same set of BSes as much as possible. A possible downside is that the computed set may not be optimal since it does not select from all candidates but only those in S_{max} . The alternative is to independently derive the set of active BSes for the idle hour. However, the computed set may not be a subset of S_{max} , thus incurring more on/off switches during idle-peak migration.

Smooth transition between idle and peak hours To minimize on/off switching and reduce energy inefficiency, we devise continuous selection for the rest of the day. It turns out that we need to switch on and switch off each BS at most once during each 24-hour duration. Figure 3.14 illustrates how our algorithm works for one grid, where

S_{max} has 11 BSes and S_{min} contains 3 BSes. The black line plots traffic profile for this grid while the dash line represents the aggregate capacity of active BSes using our scheme.

During the ramp-up transition from the idle hour with smallest traffic volume to the peak hour with heaviest traffic, we use an active node set S_t at hour t , which is always a subset of S_{max} but a superset of the previous hour S_{t-1} . That is, we find a series of active sets $S\{t\}$ that satisfying $S_{min} = S(t_i) \subseteq S(t_1) \subseteq S(t_2) \cdots \subseteq S(t_p) = S_{max}$, where $t_i < t_1 < t_2 < \cdots < t_p$ denotes the hourly sequence from the idle hour to the peak hour. The algorithmic procedure is similar to that used for the idle hour. When migrating from hour $t - 1$ to t , we only need to power on those BSes not in S_{t-1} , but retain all active BSes in S_{t-1} . If S_{t-1} is sufficient, we do not need to power on new BSes. Once a BS appears in S_{t-1} , it remains power-on at t and continues to appear in S_t until the peak hour. In our example, BSes 4-10 will switch on sequentially based on the prediction of next hourly traffic. During the other transition from the peak hour to the idle hour, the continuous selection still holds but it requires $S(t) \subseteq S(t - 1)$. In a nutshell, the Pseudo code for the graceful selection is given in Algorithm 1. $\text{SeekMinSet}(S_{prev}, S_{max}, T_t)$ is to find the set S with the minimal number of active BSes that satisfies $S_{prev} \subseteq S \subseteq S_{max}$ and the aggregate capacity of S is larger than the traffic demand T_t .

Our solution may incur suboptimal operations for energy savings when the traffic volume experiences a sudden surge (e.g., 11AM - 2PM in the example) at time t or after, before reaching its peak t_p . It keeps all current BSes on, though it may be unnecessary. We allow for this sub-optimality to minimize on/off switching. Moreover, our traces show that traffic almost follows diurnal patterns, monotonically increasing during day-time and monotonically decreasing at night. Therefore, the smooth selection works in reality by minimizing on/off switching. Our evaluation in Chapter 3.7.3 shows that the power-saving impact is no more than 1% when enabling and disabling smooth selections.

Algorithm 1 $S\{t\} = \text{SelectOnOff}(T\{t\})$

Find out the peak t_p and idle time t_i based on traffic envelop $T\{t\}$

Find out the active base station set S_{max} at t_p , as $S(t_p): S_{max} = \text{SeekMinSet}([], \text{All}, T(t_p))$

Find out the active set S_{min} at t_i , as $S(t_i): S_{min} = \text{SeekMinSet}([], S_{max}, T(t_i))$

set $S_{prev} = S_{min}$.

for each t from t_i to t_p **do**

if (no traffic load change from the previous one) **then**

 Set $S(t) = S_{prev}$

else

 Find active set between S_{prev} and S_{max}

$S(t) = \text{SeekMinSet}(S_{prev}, S_{max}, T_t)$

end if

 Set $S_{prev} = S(t)$

end for

Repeat the above loop for another period t_i to t_p in reverse time order

3.4.4 Exploit BS Diversity

Our study shows that real BS deployment is heterogenous. Capacity and coverage (i.e., maximal transmission range), as well as different power consumption models, vary at BSes. For example, a macro BS usually has larger capacity, bigger coverage and also much higher power consumption, whereas a micro BS is always less powerful. In a hybrid network, the system can save more energy by using various types of BSes during different time of the day [RFF09, RFM10].

It is known that capacity and coverage are two crucial constraints when we power off BSes to save energy. However, the bottleneck of energy saving does change at different time of the day (due to various traffic loads). For example, capacity is often bottleneck for energy saving at peak hours since many BSes have to remain on to accommodate heavy traffic load. During idle hours, coverage turns into the main bottleneck since more BSes are required to turn on to ensure full coverage while few active

BSes are unable to cover a big area.

To further improve energy savings, we propose to exploit and even inject BS diversity. The idea is to deploy dual RF subsystems at *certain* BSes to offer larger flexibility to adjust BS coverage⁵. Consequently, the BS can switch between these two systems when adjusting the coverage area at peak and idle hours. For example, for a BS in a city area, in addition to its current subsystem, we install another transceiver subsystem that works for rural areas and supports large communication coverage. Therefore, at midnight (idle hours), we enlarge its coverage by switching to the rural RF system and more neighboring BSes can turn off to reduce system-wide energy consumption. It can be illustrated by the example of Figure 3.15(c). Assume that idle traffic reduces to 10% of the peak value and each BS has the same coverage and capacity. Using the single RF operation, BS A covers at most twice its original coverage, and thus at most half of BSes can be off (Ideally⁶). On the contrary, if A is configured with another rural RF system which is able to triple its coverage or more, active BS A can allow that all its neighboring BSes power off without increasing A's capacity. This dual-mode approach reuses the existing BS components in the supporting subsystem (e.g., power supply, cooling) and most components in the communication subsystem (e.g., BBU); It requires to deploy another radio front-end (e.g., antennas and power amplifiers) only and it is cost-effective.

The BSes with higher capacity or at good locations (e.g., at the center of many micro BSes), are selected to deploy dual-RF systems. For example, macro BSes with large capacity in urban areas are good choices, Thus, these BSes switch on/off RF systems during peak/idle periods that require different transmission ranges. Note that this design is mainly applicable in the city area with dense BS deployment. Our further evaluation indicates that they do consume most energy.

While an active BS enlarges its coverage for its inactive, neighboring BSes and leads

⁵Capacity can be adjustable but we do not exploit it here because it requires radio carrier re-allocation and is more expensive in practice.

⁶In practice, constrained by BS locations, more BSes need to keep on.




Figure 3.15: Illustration of energy saving from dual-RF BS system. The dash lines represents BS coverage and red cells are ON.

to system energy savings, it may incur undesirable interference at its active, neighboring BSes. The interference degrades the performance of mobile devices associated with the victim BS. In addition, it also increases its own power consumption that might be unnecessary. One solution is to use smart antenna technique, which creates expected radiation patterns via proper beam-forming. It uses array antennas which are already available at many BSes (e.g., macro cells). In case the antenna technique is not sufficient to eliminate inference, soft frequency reuse technique can be further used to relieve Inter-BS Interference as given in [MMT08].

3.5 Comparing with the Optimal Scheme

We now analyze the performance of GreenBSN, compared with the optimal on/off scheme. GreenBSN is a distributed solution that uses a small set of active BSes based on traffic estimate. Its design eliminates several limitations of the popular optimization-driven approach, e.g., [DBM10, MCC09, BNC10]. These drawbacks include a centralized rather than distributed scheme, excessive on/off switching, unrealistic BS power consumption model, one-time optimization targeting instantaneous traffic load, and difficulty in addressing deployment diversity and node heterogeneity. Comparing with the optimal one, we next show that GreenBSN sacrifices small energy-saving gain but re-

tains *simplicity* and *practice feasibility*.

3.5.1 The Optimal Scheme

Given the runtime traffic load and BS configurations (including placement, capacity, coverage and power consumption), the optimization problem is to select the optimal scheme S^* , which minimizes energy consumption while satisfying the constraints on coverage and capacity,

$$S^* = \arg \min_S \sum_i P_i(S, L), \quad (3.7)$$

where $P_i(S, L)$ is the energy consumption of BS i using scheme S under traffic load L . The scheme S determines the set of active BSes as well as traffic allocation. The problem formulation is similar to the user-cell association in [DBM10, FBG11].

We first use a simple example to show why the optimization-based scheme may outperform GreenBSN in certain scenarios. Figure 3.16(a) plots a deployment map of nine BSes in Region 1. Four virtual grids are formed: $\{1, 6, 8\}$, $\{2, 5\}$, $\{3, 4, 7\}$ and $\{9\}$, following our “northeast” rule. Figure 3.16(b) marks the active BSes for each hour (the blue ‘+’ denotes our grid algorithm and the red circle is the optimal one). The optimal one is obtained via an unrealistic, exhaustive search. It is seen that the optimal one powers on fewer BSes and reduces more energy. The reason is that our grid-based algorithm requires at least one active BS for each grid even though it may not be necessary (e.g., when the grid can be covered by BSes in neighboring grids). For example, the *grid* scheme has to turn on BS 5 (at midnight) for the grid $\{2, 5\}$ whereas the *optimal* one leverages BSes 6 and 8 to cover BS 5, and BSes 3 and 4 to cover BS 2. Consequently, there is no need to turn on BSes 2 and 5 under light traffic. Moreover, capacity may not be fully utilized due to lack of coordination among neighboring grids. For example, due to heavy traffic at noon, grid $\{3, 4, 7\}$ has to power on two BSes (4 and 7), and grid $\{1, 6, 8\}$ turns on two BSes (1 and 8). In contrast, the *optimal* scheme leverages the extra capacity from neighboring grid $\{6, 8\}$, thus powering off BS 4, as

shown in Figure 3.16(b).

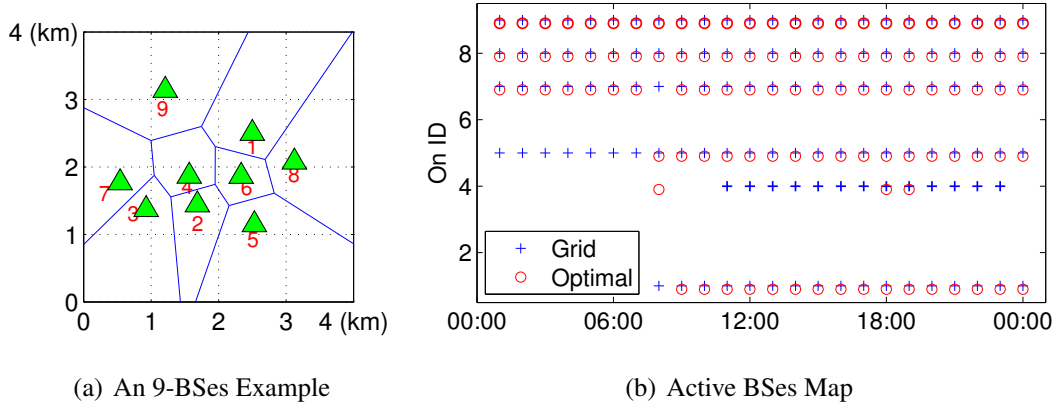


Figure 3.16: Grid-based versus optimization-based schemes.

3.5.2 Bounded performance gap

The optimization scheme is able to exploit cross-grid deployment redundancy to improve its energy-saving gain. However, this optimization problem is NP hard. We can readily show that, when each BS has identical power model, the performance gap between the grid-based scheme and the optimal one is upper bounded, as stated by Theorem 3.5.2 and Corollary 3.1. Now we elaborate on the details.

Lemma 3.1. *In case of homogeneous power models, the optimal scheme is equivalent to the one with minimum number of active BSes that accommodates the aggregate network traffic.*

Proof. Let the power model be $P(L) = \omega L + \theta$ and $|S|$ be the number of active BSes in the scheme S . Thus, the energy consumption for this scheme is

$$E = \sum_i P_i(S, L) = \sum_i (\omega L_i + \theta) = \omega \cdot \sum L_i + |S| \cdot \theta. \quad (3.8)$$

The first part is constant under given traffic. The optimal solution to minimizing energy consumption is hence equivalent to the one with smallest number of active BSes. \square

Theorem 3.5.1. *The optimization problem to minimize energy consumption is NP-hard.*

Proof. We construct one problem instance with homogeneous power models. Based on Lemma 3.1, energy minimization is to find the minimum number of active BSes. This is equivalent to identifying *the minimum dominating set (MDS)* in the corresponding graph where vertices and edges denote the BSes and their equivalency relationship, respectively. Therefore, our optimization problem is no easier than the *MDS* problem (a special case) that is known to be NP-hard. \square

Theorem 3.5.2. *In case of homogeneous settings with identical power model and capacity at each BS, our solution has at most $(q - 1)$ more active BSes than the optimum, where q is the number of grids in the network.*

Proof. Let $n(i)$ and $n^*(i)$ be the number of active BSes in grid-cell i from our solution and the optimum, respectively. Let C be the capacity of each BS and $L(i)$ be the aggregate traffic demand in cell i . Since our grid scheme selects the minimum number of active BSes that provides capacity larger than traffic load in each cell, we have

$$n(i) \times C \geq L(i) > (n(i) - 1) \times C, \quad \forall i. \quad (3.9)$$

On the other hand, the optimal solution should provide the total capacity larger than the total traffic load:

$$\sum_i n^*(i) \times C \geq \sum_i L(i) > \sum_i (n(i) - 1) \times C. \quad (3.10)$$

Therefore, we get:

$$\sum_i n^*(i) > \sum_i (n(i) - 1) = \sum_i n(i) - q, \quad (3.11)$$

where q is the number of grids in the network. \square

It is natural to extend Theorem 3.5.2 to a more general case with heterogeneous capacity settings.

Corollary 3.1. *Let η be the maximum capacity ratio between any BS i and other BSes covered by BS i , $\eta = \max_{\forall i,j} C_i/C_j$ where BS j can be covered by BS i . In case of homogeneous power models, our solution has at most $\eta n^* + (q - 1)$ active BSes where n^* is the optimal number of active BSes.*

Proof. Let $S(i)$ and $S^*(i)$ be the set of active BSes in grid-cell i for our solution and the optimum, respectively. Let $I(i) = S(i) \cap S^*(i)$, $D(i) = S(i) - S^*(i)$, and $D^*(i) = S^*(i) - S(i)$. In the *grid* scheme, the traffic in cell i is served by BSes in $I(i)$ and $D(i)$. In contrast, for the optimal scheme, the traffic in cell i is served by BSes in $I(i)$, $D^*(i)$ and $O^*(i)$, where $O^*(i)$ denotes the set of BSes that are not in cell i but can serve cell i . Consequently, in the *grid* scheme, the traffic load $LD(i)$ served by the BSes in $D(i)$ is larger than the aggregate capacity of all $|D(i)| - 1$ BSes; otherwise, it is unnecessary to power on all these BSes. On the other hand, the *optimal* scheme also needs to offer capacity larger than this load. We use $Cap(S) = \sum_{j \in S} C_j$ to denote the capacity provided by the BS set S . We then have

$$Cap(\{D(i)/\epsilon \in D(i)\}) < LD(i) \leq Cap(\{D^*(i)|O^*(i)\}), \quad (3.12)$$

$$\sum_i Cap(\{D(i)/\epsilon \in D(i)\}) < \sum_i Cap(\{D^*(i)|O^*(i)\}). \quad (3.13)$$

It is known that the capacity of the BS in the $D(i)$ is at least $1/\eta$ of the one in $D^*(i) \cup O^*(i)$. Therefore, we have

$$\sum_i (|D(i)| - 1) < \eta \cdot |\cup_i (\{D^*(i)|O^*(i)\})|. \quad (3.14)$$

Let x and x^* be the sum of $|D(i)|$ and $|D^*(i)|$, respectively. Using $n - n^* = x - x^*$, we can show that $x - q < \eta \cdot (n^* - (n - x))$. We then have $n - n^* < x - \frac{x-q}{\eta} < n - \frac{n-q}{\eta}$. It follows that $n < \eta n^* + q$. \square

In reality, BSes in close proximity are probably similar, e.g., using the same type of BSes in one area. Therefore, η is probably not so big that the performance degradation in our scheme is reasonably bounded in practice. We note that power models may vary at BSes (i.e., having different ω, θ due to configurations). However, the gap would not be significant; otherwise, more energy-efficient ones would be deployed.

On the other hand, the optimization-based scheme mainly exploits local deployment redundancy to improve its energy-saving gain. It may power off more BSes only if each doze-off BS can be covered by several active BSes. However, this redundancy is

ultimately limited by the local deployment density, which is always bounded in reality. Therefore, the optimization scheme cannot yield larger gain as the area further increases (shown in Figure 3.20). The fundamental reason is that, energy savings are ultimately decided by node deployment density, capacity and coverage. No matter what scheme we use, the selected BSes only work within their deployment and coverage proximity. As long as the deployment density is bounded, the gap between different schemes is also bounded.

3.5.3 Tradeoffs for practicality

Compared with the optimal scheme, our profiling and smooth ON/OFF scheme may also incur reduced energy savings. In this work, we trade optimal power savings for solution simplicity. We can show that this reduced energy-saving gain is marginal. The reduced energy saving due to profiling only happens when our profiles overestimate traffic load and this overestimation triggers to change the on/off scheme (e.g., turning on more BSes). Fortunately, our study shows that the traffic profile offers an accurate traffic upper bound estimate, thus leaving little room for capacity waste; Moreover, not all overestimates lead to more active BSes since each BS capacity is usually discrete. Each BS capacity may have spare room to accommodate the overestimate incurred by the profiling scheme, thus avoiding more BS activations. The reduced energy saving due to smooth ON/OFF selection might be caused by two factors. First, the daily traffic does not follow diurnal patterns, monotonically increasing during day-time and monotonically decreasing at night. For example, the traffic drops at one slot when it goes up in the following slots. It may waste slight energy only if some active BSes that power on early become unnecessary at that time. Second, the optimal scheme could be different from the one using BSes from the specific BS set (e.g., S_{max} for peak hours). Our study shows that both cases happen rarely and the gap due to smooth ON/OFF scheme is no more than 1% in Chapter 3.7.3.

3.6 Working within 3G Standards

GreenBSN has to be standard compliant and address practical issues: (1) How to let active BSes cover the communication area of those sleep ones; (2) How to migrate existing user clients from the about-to-sleep BSes to other active BSes; (3) How to leverage the 3G infrastructure to share traffic information among local BSes in a grid; (4) How to coordinate the operations of cooling subsystems and NodeB communication subsystems during the energy-saving process; (5) How to handle unexpected traffic surge.

Figure 3.17 plots the overall GreenBSN work-flow for BSes and RNC. RNC collects traffic information from all the BSes and generates the on/off map which coordinates the switching on/off of all the BSes in one grid. Upon on/off signals, active BSes (i.e., onBS) and sleep BSes (i.e., offBS) take appropriate actions to enable on/off schedule, for example, the onBS uses cell breathing to expand its coverage when neighboring BSes power off. We now elaborate on the detailed implementation.




Figure 3.17: Workflow of GreenBSN solution.

Adjusting BS Coverage via Cell Breathing In our scheme, some BSs need to extend their coverage to serve clients originally covered by neighboring BSs that will

power off. To this end, we leverage the well-known “cell breathing” technique that adjusts cell boundaries in today’s 3G networks [3GP08c, NWG10]. Cell breathing is traditionally used to adjust the cell size based on the number of client requests to achieve load balancing or capacity increase through micro-cell splitting. We use it for the alternative purpose of power savings. Specifically, the effective service area expands and contracts according to the energy-saving requirement. By increasing the cell radius, an active BS can effectively extend the coverage area to neighboring BSs. Note that most Node B vendors offer products operating over a wide communication range (say, 200m to a few kilometers).

User Migration by Leveraging Handoff When migrating users from the original BS to the equivalent BS for power savings, we leverage the network-controlled handoff (NCHO) mechanism in 3G standards currently used for mobility support. Figure 3.18 shows the migration procedure of mobile users to the other active BS when the serving BS decides to power off. For each active UE in the original BS (OBS), the following procedure is performed: (1) The OBS sends a handoff request to the neighboring active BS (ABS) via RNC; (2) The ABS acknowledges the handoff request and reserves resources for the migrating UE; (3) Upon receiving the handoff ACK from the ABS, the OBS sends the UE a handoff command; (4) The UE executes the handoff command via new association with the ABS. Then this handoff process is done in NCHO [3GP10]. In case of handoff failures, the OBS may repeat the above procedure with other active BSes until all UE handoffs succeed or time out. The OBS will defer its power-off if some UEs are still associated with it. Note that our handoff triggering event (i.e., BS power on/off) does not require additional modification to the current 3G NCHO operation except adding one more event type. Note that this handoff process is applied to each active UE in the OBS. To speed up the UE migration, we can enable batch handoff requests for multiple users. The migration process can thus be readily made 3G standard compliant.

Information Sharing via RNC To coordinate BS operations within the same grid,




Figure 3.18: 3G Handoff procedure for user migration.

BSes in the grid should exchange information, for example, traffic information to compute the envelope for the aggregate traffic, and on/off BS status to guide cell breathing and handoff. A natural place to exchange such information is via the RNCs. In typical cases, BSes belonging to the same grid also own the same RNC, which is the natural hub for such information exchange and aggregation. In the extreme case that BSes in a given grid do not belong to the same RNC, we can modify the grid construction procedure by imposing the condition that only equivalent BSes belonging to the same RNC can form a grid. The downside of this change is that it may increase the number of grid cells, but with the benefit of reducing inter-RNC message exchange.

Coordinated operations within NodeB Most Node B subsystems require proper operating temperature to function well. When powering off the entire BS for an extended period of time, the ambient temperature may change beyond the desired operating value. Therefore, before powering on the Node B subsystem, we need to power on the cooling/heating subsystem in advance. Our measurements done at three real-life BS machine rooms show that 30 minutes are generally enough for the current cooling/heating system to bring the room temperature to the desired value.

Emergency BS power-on While our profile-based approach typically gives a reliable estimate on the traffic envelope, rare-case traffic surge may also occur. To prepare for

such transient surges, each active BS monitors its traffic volume. Whenever it sees sudden surge beyond the envelope specified by the profile, it will notify its RNC. The RNC will subsequently trigger emergency power-on for the neighboring power-off BSes. The number of power-on BSes depends on the traffic surge volume the RNC is notified.

In a nutshell, GreenBSN is 3G standard compliant; By reusing the existing cell breathing, handoff and RNC control function, it is implementable in the existing BS networks. Algorithm 2 shows the processing to update on/off BSes at RNC. Moreover, its three key components of BS clustering, traffic profiling and graceful BS selection, can be performed mostly offline which incurs low runtime overhead. Thus GreenBSN is realistic and readily applicable to the current BS system.

Algorithm 2 calcOnOffBS() //at each RNC

- 1: Obtain virtual grids G and initial traffic envelop EV as given in Sect. 3.4.1 and 3.4.2.
- 2: Calc. active set $S(t)$ for each grid as given in Sect. 3.4.3
- 3: **while** (1) { //online processing loop, at the t-th hour
- 4: **for** each grid G_i **do**
- 5: Request real-time BS traffic and sum up $\rightarrow L_i(t)$
- 6: UpdateProfile() and obtain EV' as the updated envelop
- 7: **if** $EV'_i(t) > \xi \sum_{j \in S_i(t)} C_j$ ($\xi = 1$ - reservation factor) **then**
- 8: update $S'_i(t) = S_i(t) + \Delta$ s.t. $\xi \sum_{j \in S'_i(t)} C_j \geq EV'_i(t)$
- 9: **end if**
- 10: **if** $L_i(t) > EV'_i(t) \ \&\& \ L_i(t) > \xi \sum_{j \in S'_i(t)} C_j$ **then**
- 11: update $S'_i(t)$ s.t $\xi capacity \geq L_i(t)$ // emergence BS on
- 12: **end if**
- 13: $S_i(t) = S'_i(t)$; The on-set is $S_i(t) - S_i(t-1)$; the off-set is $S_i(t-1) - S_i(t)$ //One of the
On and Off set is empty
- 14: **end for**
- 15: }

3.7 Evaluation

We evaluate our power-saving solution using two-month traffic traces collected from four regional 3G networks. We use the first five-week data to construct traffic profiles, and use the remaining three-week traces to assess our solution. We do not use traces from the second operator since the traffic record is too short to create traffic profiles.

Evaluation setting We first evaluate our solution in default parameter settings: (i) profiling parameter $\gamma = 3$; (ii) heterogeneous BS capacity being 110% of the maximum traffic load at a given BS; (iii) power model $P_{tx} = 6L + 600w$ and $P_{misc} = 1500w$ at normal transmission range; $P_{tx} = 12L + 600w$ when expanding to the maximum transmission range. This model states that consumed power still grows linearly with the load but with a larger coefficient, say, P_a doubles at maximum coverage; (iv) the maximal transmission range of 1–2 km, consistent with many available products. We also gauge the effect of various parameters and other power model alternatives, and compare our solution with the optimization-based approach later.

3.7.1 Overall Performance

Table 3.6 summarizes the results on the above default setting. The table presents the total daily energy consumption of the current 3G network E_{old} , BBU-standby solution E_{bbu} , and our power saving scheme E_{our} , the average energy-saving percentage, and the daily miss traffic (due to profiling inaccuracy or capacity limit) and the active BS count using our scheme. The BBU solution, proposed by some BS vendors [Eri08], aims to turn off some sub-carriers and place BBU into the standby mode when the traffic load is low. We also separate weekday and weekend performances, but they are similar. We tried to compare with other energy saving scheme in [MCC09, DBM10, BNC10, SKY11, FBG11, RFF09, BOL09, NWG10], but most theoretical work could not be applied due to their assumption constraints; the hybrid deployment tech-

nique in [RFF09, BOL09] are orthogonal to GreenBSN. The cell breathing technique in [NWG10] is already used in our scheme.

	Region 1	Region 2	Region 3	Region 4
E_{old} (Mwh)	9.81(9.70/9.86)	2.63(2.52/2.67)	8.58(8.45/8.64)	9.18(8.98/9.29)
Daily E/node (kwh)	55.4(54.7/55.7)	58.4(56.0/59.3)	55.7(54.9/56.1)	56.0(54.5/56.6)
E_{bbu} (Mwh)	8.3	2.3	7.5	7.9
E Gain	15.7%	13.8%	12.6%	13.9%
E_{our} (Mwh)	4.64	1.40	5.94	7.03
E Gain(%)	52.7%	46.6%	30.8%	23.4%
(min–max)	(34.2–75.9)	(20.6–76.1)	(16.5–46.6)	(9.9–35.4)
#miss/BS	2.83	5.23	4.37	0.12
missRatio(%)	6.7e-4	7.9e-4	8.16e-4	1.86e-5
#BS(weekday)	34–97	8–32	79–122	104–142
#BS(weekend)	34–85	8–19	79–107	103–122

Table 3.6: Power saving in four regions.

	Daytime	Midnight	A(sparse)	B (dense)
Region 1	40.7%	73.7%	28.1%	61.6%
Region 2	31.2%	71.6%	27.7%	55.3%
Region 3	20.9%	45.6%	8.8%	51.3%
Region 4	15.6%	34.7%	7.9%	30.8%

Table 3.7: Power saving in peak/idle hours and subregions.

We make four observations. *First, significant power saving is feasible.* Our profile-based scheme achieves average daily energy savings about 50% in Regions 1 and 2 (dense areas) and 20–30% in Regions 3 and 4 (sparse areas). Compared with the BBU-standby solution, our scheme yields more power savings because the BBU scheme saves P_β but cannot eliminate P_{misc} . In all cases, 15%–40% BSes are powered on/off in the regional network, i.e., 20–60 BSes each day. Those BSes switch on/off *only once* during each 24-hour period, confirming the operational simplicity of our scheme.

Second, the power-saving gain is mainly attributed to traffic diversity and deployment density. Since the wasted energy is unproportionally large under light traffic, our scheme achieves the largest energy savings during idle period. Table 3.7 shows that, the power-saving gain reaches as high as 70% during night time in Regions 1 and 2, and the gain at night time is about 2x the value at daytime in all regions. Deployment density is another crucial factor to power savings. It determines the degree of redundancy to turn off BSes. The gain in dense areas can reach 30–60%, almost 2-6x the value in sparse areas. It also explains why the gain is lower (23.4%) in Region 4, while reaching 52.7% in Region 1. We also assess the impact of grid formation. We find that, the gain is similar no matter what grids are used with different BS sets. The location-dependent density is the key factor of power savings, and grid formation does not affect the inherent density.

Third, we can save energy by powering off some BSes even during daytime, particularly in dense areas (e.g., Region 1). Our analysis shows that traffic multiplexing over time and in space is the main contributing factor to such a gain. The current BS deployment does not take the broad system view, but seeks to meet the peak traffic requirement at each location myopically. Consequently, it is inevitable for the operator to over-provision the capacity too aggressively, as observed in Figures 3.7(a) and 3.7(b).

Last, our results also reveal the tradeoff between power savings and performance degradation. Due to occasional unavailability of spare capacity, some user requests will be denied. In principle, our profiling scheme may not always capture extreme-case traffic surge in its profile envelope, thus leading to transient overload beyond the provisioned capacity. However, our study shows that such cases are rare. The average miss ratio is kept as low as $< 0.1\%$, i.e., up to 6 requests per BS each day. If needed, we can use more conservative policies (e.g., use larger profiling margin γ) or leverage the emergency BS power-on mechanism, to further reduce the miss ratio.

3.7.2 Impact of Various Components

We now study the effect of various components and parameter settings on our energy-saving performance.

BS power models The transmission power P_{tx} and the cooling power P_{misc} vary with the sector count and with seasonal changes, respectively. Different vendor products also introduce diversity into power models. Table 3.8 presents the assessment of six power models given in Table 3.1. The first five models are homogeneous and the last one assesses the tradeoff between high capacity and high energy efficiency, where BSes with larger capacity consume more power. Because results are similar in other regions, we only show a case in Region 1.

			Region 1		
#	Model	Setting	E_{old}	E_{our}	E gain
1	Winter	$P=1000+6L+600$	7.7K	3.6K	53.2%
2	Summer	$P=2000+6L+600$	11.9K	5.3K	55.4%
3	Sp/Fall	$P=1500+6L+600$	9.8K	4.6K	53.1%
4	Sp/F-low	$P = 1500+4L+600$	9.5K	4.2K	55.8%
5	Sp/F-high	$P = 1500+8L+600$	10.3K	5.0K	51.4%
6	Hybrid	H, $P = 2000+8L+800$ M, $P = 1500+6L+600$ L, $P = 1000+4L+400$	8.8K	5.00K	43.2%

Table 3.8: Energy saving with different power models.

Our results show that the power model diversity does not much affect the saving gain. It only leads to visible changes in the absolute energy consumption. The power-saving percentage is almost invariant in the five homogeneous cases. It drops about 4-8% in all regions in Model 6, caused by the tradeoff between energy efficiency and capacity. Power models do not affect the miss rate and active sets. The stable power-saving gain in different models indicates that our scheme can work well around the

whole year.

Profiling scheme We assess the profiling parameter γ by varying $\gamma = 1, 2, \dots, 5$. We also compare the grid-group profiling and the individual BS profiling scheme. Figures 3.19(a), 3.19(b), and 3.19(c) plot power-saving gain (with min/max bounds), miss requests per BS, and the active BS percentage, respectively. The larger γ tends to over-estimate the traffic load and turn on more active BSes. Consequently, the set of active BSes $|S_{max}|$ grows larger and the energy-saving percentage is reduced. The results show that, when γ grows up from 1 to 5, daily energy-saving gain only decreases about 5–10%, which offers large freedom to set γ . On the other hand, the number of miss requests per BS decreases significantly. The larger γ tends to over-estimate the traffic load. We also see that group profiling outperforms the individual one on energy-saving gain and miss rate. Such a gain can be attributed to the effect of traffic aggregation in a local proximity and multiplexing over time, as shown in Figure 3.7(a).

BS capacity We vary the BS capacity by multiplying a variable η (from 1 to 2) and the peak traffic load; we also use heterogeneous BS capacity as shown in Figure 3.11. The ratio of aggregate capacity to aggregate peak traffic is about 1.8. Figures 3.19(d), 3.19(e), and 3.19(f) plot power-saving gain, miss requests per BS, and the active BS percentage, respectively. With a larger BS capacity, the network reduces the active BS count while serving the traffic demand, thus reducing power waste. When the BS capacity doubles, energy-saving gain increases by about 10%. However, as we vary BS capacity, the maximum energy-saving gain (when traffic is lightest) is almost invariant. It implies that energy saving under light traffic is constrained by coverage rather than capacity. On the other hand, the minimum energy-saving gain (when traffic is heaviest) increases significantly (e.g. 2x-4x in Region 2, 3, and 4) when the capacity doubles. This case of power saving is mainly constrained by capacity. We also find that increasing BS capacity can offset the profiling inaccuracy, while larger capacity may trigger more BSes to power off, leading to higher miss rate, though there is an exception that the miss ratio at $\eta = 1.5$ is larger than the one at $\eta = 1.2$ in Region 3. It is because both







Figure 3.19: Evaluation results of various effects on energy-saving, miss traffic, and the number of active BSes.

active BS sets are selected to afford the traffic demand given in the profile; However, the aggregate capacity at $\eta = 1.5$ might be smaller than the 1.2-one when fewer BSes are on. In the extreme case that the aggregate capacity is not sufficient, it incurs more misses. Note that, in all cases, the absolute number of miss requests per BS remains small (<10 , with miss ratio smaller than 0.2%).

BS maximum transmission range We vary the maximum transmission range of a BS by multiplying variable κ and the normal transmission range; we vary $\kappa = 1, 2, 3, 5$ in our study. When $\kappa = 1$, we set $R_i = \max(500, r_i)$. We also compare them with homogeneous settings of $R = 1\text{km}$ or 2km based on BS deployment analysis of [Mis04]. Figures 3.19(g), 3.19(h), and 3.19(i) plot power-saving gain, miss requests per BS, and the active BS percentage, respectively. In general, the larger the transmission range, the more active BSes the network can reduce to cover the entire area. Our results show that we achieve significant power savings when the maximum transmission range is three times larger than the operational range. When it is very small, coverage becomes the limiting factor.

Other design variants We also evaluate two more design variants in our scheme. The first is to power on the sleeping BS ahead of the expected working time. It is to give enough time for the cooling system to adjust the ambient temperature inside the BS. The second option is to always reserve a fraction (say, 10%) of the capacity in a BS to be prepared for the worst-case scenario (e.g., unexpected or transient traffic surge). Figures 3.19(j), 3.19(k), and 3.19(l) plot power-saving gain, miss requests per BS, and the active BS percentage, respectively. The results show that, both design variants have little impact on energy saving. The first option (ahead of time) decreases only 1–2% in energy-saving gain, while the 80% resource reservation only reduces 5-10% energy saving gain. Neither visibly affects the miss rate.

3.7.3 Comparing with the Optimization-based Scheme

We now compare our solution with the optimization-based scheme. We compare each of the three design components, i.e., *virtual grid*, *profiling*, and *graceful selection*, with the corresponding idealized or optimization-based solution.

Virtual grid Our grid-based scheme decouples location-dependent coverage and capacity constraints. The BSes are divided into virtual grids so that BSes in a grid can cover each other, and each grid then makes decision independently based on capacity requirement. Given the same traffic load, we now compare it with the optimization-based approach, which uses brute-force search to select active BSes that consume minimum energy while satisfying both capacity and coverage constraints. The exhaustive search offers an upper bound for energy savings, even compared with other proposed optimization-based solutions in [DBM10, FBG11].

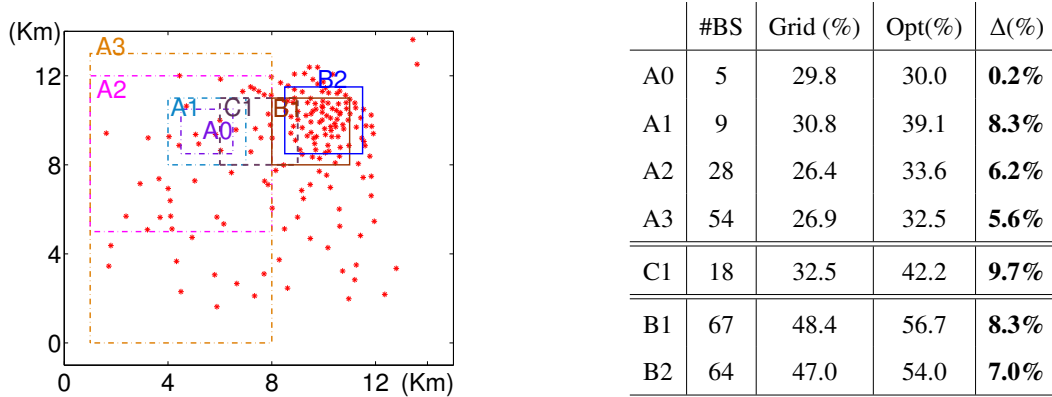


Figure 3.20: Comparison with the optimization-based scheme in different cases. Left: case settings; Right: energy saving.

We compare both schemes in sub-regions of different area size and different deployment density in Region 1: $A_0 - A_3$ for sparse deployment, $B_1 - B_2$ for dense deployment, A_1, C_1, B_1 have same area size ($3\text{km} \times 3\text{km}$) but with different deployment density, as shown in the left of Figure 3.20. The right table lists the energy-saving gains by both schemes in each sub-region. These results show that the energy-saving gap between our scheme and the optimization scheme is not large, less than 10% in all the

cases. We also make interesting observations. When the area size is small, the performance gap also tends to be small. As the area size gradually increases, the optimization scheme has larger space to optimize, thus exhibiting larger performance gap. However, when the area size further grows, the gap saturates since the deployment density now becomes the limiting factor. Compared with our scheme, the optimization scheme can exploit cross-grid deployment redundancy but this redundancy is bounded by inherent deployment and coverage proximity. Therefore, the optimization scheme cannot yield more gain as the area further increases.

The gap between our scheme and the optimization scheme is bound mainly exploits the local deployment redundancy to improve its power-saving gain. It may turn off more BSes only if each doze-off BS can be covered by several active ones. But this redundancy is ultimately decided by the local deployment density, which is always bounded in reality. Therefore, the optimization scheme cannot yield larger gain as the area increases further. The fundamental reason is that, energy savings are ultimately decided by node deployment density, capacity and coverage. No matter what scheme we use, the selected BSes only work with their inherent deployment and coverage proximity. As long as the deployment density is bounded, the gap between different schemes is also bounded.

Profiling We use traffic profiles (i.e., estimated traffic envelope) rather than runtime traffic to guide our BS activation and deactivation. It is inevitable to overestimate the runtime traffic, and tends to turn on more BSes occasionally. We compare the energy-saving percentage using our scheme based on traffic profiles and the one based on runtime traffic. The results are plotted in Figure 3.21(a). Our scheme yields 52.7%, 46.6%, 30.8%, and 23.4% of energy savings in four regions, respectively. It shows that the optimization based on the runtime traffic yields 57.1%, 49.8%, 35.3%, and 26.0% energy saving in Region 1-4, respectively. That is, the reduced energy saving caused by the profiling scheme is small (<4.5%). This shows that our traffic profile estimate is simple and accurate.




Figure 3.21: Comparison of our scheme and other designs on energy saving: (a) using real-time traffic; (b) disabling smooth selection.

As explained in Chapter 3.5, two factors help to achieve energy savings in our scheme: the accurate traffic profile estimates and tolerance to traffic estimation error that discrete BS capacity provides. Our traffic profile estimation offers an accurate traffic upper bound estimate, thus leaving little room for capacity waste; Second, not all the overestimate lead to more active BSes. Each BS capacity may have spare room to accommodate the overestimate incurred by the profiling scheme, thus avoiding more BS activations.

Graceful selection We also compare our graceful selection with the idealized option that disables smooth selection and maximizes energy saving for each time slot. To reduce the frequency of ON/OFF switches, our scheme selects the active BSes from those already active ones. However, our scheme may lead to retaining unnecessary BSes active or not selecting the most energy-efficient BSes. We compare the energy savings using our scheme (via smooth selection) and the option that disables smooth selection. Figure 3.21(b) shows the energy-saving percentage of both schemes in four regions. The scheme without smooth selection yields energy-saving gains of 52.8%, 47.2%, 31.2%, and 23.4%, in Regions 1-4, respectively, leading to 0.1–0.7% differences in saving gains compared with our scheme. The energy-saving reduction due to smooth selection is thus negligible (<1%). It shows that the energy-saving reduction

due to graceful selection is negligible in practice. The reason is that, the traffic envelope may be not monotonically increasing only in rare cases. Therefore, the chance is slim when switching off some BSes that are already on for a short period of time but again powering them on later. It implies, our smooth selection has little negative impact on energy savings, while ensuring the simplicity of smooth BS ON/OFF operations.

3.7.4 Impact on Clients

Our scheme does pay a cost to achieve energy savings on the infrastructure side. It will increase transmit power at client devices when sending uplink data traffic during idle hours (say, late evenings or weekends). In our scheme, when the closest BS powers off, a mobile client will migrate to an active but distant BS, thus incurring additional energy for *uplink transmissions*. However, its impact on the client device is not as severe as it appears. First, the uplink traffic volume is far less than the downlink, which is dominant. The uplink-to-downlink traffic ratio is about 1:8 in the Internet, and 3G networks have similar ratios observed from our traces and prior measurement [FLM10]. Second, the transmission range-extended BSes are conceptually equivalent to the BSes deployed in rural areas, which have larger coverage. Current client devices do not seem to experience severe energy penalty in rural areas. Finally, mobile users are more likely to be uniformly distributed around their serving BSes. Therefore, only a fraction of users will increase their power for uplink transmissions, as shown next.

Our scheme may have other negative impacts on mobile clients. When on/off schedule takes effects, mobile clients may be mitigated to a further but active BS. It may degrade user performance (throughput, latency etc) and increase energy consumption for *uplink transmission* on client side. It may affect the wireless channel quality, and thus degrade user performance (throughput, latency etc) and increase energy consumption on client side. Based on our domain knowledge, we can assume that the change of transmission change is the major factor.

	Region 1		Region 2		Region 3		Region 4	
	70th	90th	70th	90th	70th	90th	70th	90th
4 AM	588	920	728	991	310	823	0	749
10AM	64	329	46	159	0	329	0	281
4 PM	0	297	0	0	0	17	0	0
10PM	92	401	176	486	0	341	0	600

Table 3.9: The BS-to-client distance change (m).

We quantify the change in transmission range due to our power-saving scheme. Assume uniform distributions for users in their original BSes. Users also associate with their closest active BS. Figure 3.22 plots the BS-to-client distance change over time in two regions. We see that during daytime, more BSes are active and the distance change is negligible. At night, the distance may increase, e.g., up to 1 Km for 10% clients in Region 1. Table 3.9 shows the distance change for n -th user at a specific time in four regions; it shows that the affected number of clients is still under control.




Figure 3.22: Transmission range change in Regions 1 and 4.

3.7.5 Evaluation Summary

The real trace-based evaluation validates our power-saving solution and shows that significant energy-saving is feasible. It yields up to 52.7% savings in a dense area, and 23.4% in a sparse area. Savings are more significant during night time, e.g., up to 70%,

and 1 or 2x larger than those during daytime in all regions. Even during daytime, 20–40% savings are possible by exploiting temporal-spatial multiplexing gains. The traffic miss ratio is also kept lower than 0.1% in the worst case with having the appropriate number of BSes (15–40%) switch on/off at most once during each 24-hour period. Evaluations on various parameters also confirm that our solution is readily applicable to various practical scenarios. For the tradeoff between power-saving gains and the miss rate, our scheme can achieve high power savings as well as low miss rate, e.g., less than 0.1% in our solution. Compared with the optimization-based exhaustive search, our solution achieves effective power savings in a simple and practical manner, while keeping the gap less than 10% in all tested scenarios. On the downside, our scheme may incur increased energy consumption on the client side, but only for uplink traffic and mostly during light-traffic night time.

CHAPTER 4

Mobile Data Accounting

In this chapter, we move to the second topic of mobile data accounting. We first address the problems in mobile data accounting and introduce our study methodology. We then disclose the limitations in mobile data accounting architecture and policy and present unexpected accounting behaviors in various cases. For each limitation, we also identify the root cause and propose remedies.

4.1 Problems in Mobile Data Accounting

We start with a motivative example and then identify the issues to address.

4.1.1 A Motivative Example

Unlike the flat charging practice over the Internet, the current 3G network has been using usage-based charging for its data services. A mobile user is charged of a monthly bill based on his/her used data volume or the cap. Undoubtedly, mobile data accounting, which records how much data volume is *actually* used, plays a critical role.

The current mobile data accounting has been operational for a few years and the

practice has been generally successful. However, this system mostly works as a black box for users, and users do have questions and concerns. Consider Alice, a typical 3G user, as an example. Alice just received a monthly bill of \$25.8 for 387.4MB data usage, with a portion of the bill being shown in Table 4.1. Even with this itemized data usage, Alice may still have lots of doubts and questions in her mind:

1. Does my iPhone really use 4385KB (but not 2.3MB less) since I remembered I only downloaded a 2MB app from the App Store?
2. For the 31KB item, I remembered I clicked an invalid web page link that did not show me any real content. Why should I be charged?
3. The item at 4:50AM is even ridiculous. Definitely, I was asleep then and did nothing with my iPhone. Why did a 16MB item show up in my bill?
4. Moreover, how can I find out if the operator made a mistake and overcharged me? Anyway, I heard that up to 20M Americans using their iPhones/iPads are over-charged by 7-14% on average and up to 300% in some cases [Rep11].
5. If the operator indeed made a mistake, how much did I overpay?
6. Is there any chance I can somewhat evade the charging system and pay less?

The list goes on and long. From the consumer perspective, Alice has every right to know the answers since it is about the money matter. On the technology side, answers to all these questions reside on the accounting system used by 3G/4G networks.

4.1.2 Issues in Data Accounting

Our goal is to address these user concerns. We aim to investigate mobile accounting for data services and examine how operator policies and their enforcement affect our pockets. The research focuses on accounting, which records the usage volume over time

Date	Time	Type	Direction	Msg/Kb/Min
03/04	09:21PM	Internet/Media Net	Sent	31KB
03/04	08:07PM	Internet/Media Net	Sent	4385KB
03/04	06:16PM	Internet/Media Net	Sent	65KB
03/04	02:45PM	Media Messaging	Sent	12KB
...
03:04	04:50AM	Internet/Media Net	Sent	16782KB

Table 4.1: Example of itemized mobile data usage.

for each user, rather than pricing that sets the unit price for usage (e.g. whether the unit-price, e.g., 10cent/MB, is too expensive). The latter problem is driven by marketing and cost factors, beyond our scope and interest. To this end, we examine the 3G accounting architecture, as well as the operators' policy practice. Little work has been done so far to understand how effective both work in reality.

We study the implication of such architectures and policies from a user-centric perspective. Our evaluation is to compare the user-recorded data volume with the network-recorded usage. Our study is user centric overall, in that *users pay for what they actually get at the end systems*. Anyway, the end systems are where users obtain their data service. This user-centric guideline may not necessarily concur with the infrastructure-oriented view taken by operators. When conflicts happen, one camp may suffer. In this work, our main goal is not to take specific position on what side to stand with, but to illustrate and quantify the discrepancy in different scenarios. We collect the actual usage at the two end systems, and compare with the volume given by the 3G accounting system. The goal is to identify possible limitations and existing loopholes though such cases may occur rarely in reality, and demonstrate their effect on end users. We explore two dimensions of the problem from both robustness and security perspective.

- *What*: What is to be accounted? What is the difference in accounting/charging for different types of data traffic?
- *How*: How does 3G accounting handle various cases of end-to-end data delivery?

	<i>What</i>	<i>How</i>
<i>Robustness</i>	free-data service (Chapter 4.3) App signaling (Chapter 4.6)	Accounting under wireless loss (Chapter 4.4) Accounting with middleboxes (Chapter 4.6)
<i>Security</i>	free-data attack (Chapter 4.3) stealthy-spam-attack (Chapter 4.5)	stealthy-spam-attack (Chapter 4.5)

Table 4.2: Dimensions of mobile accounting issues

- *Robustness*: How the system works in reasonable scenarios where all the operations are expected and no malicious exploits are performed?
- *Security*: What if the malicious users launch specific attacks against the system?

The “*what*” issue depends on the charging/accounting policy. Each 3G operator can define its own application-specific policy on charging. Along this line, we are particularly interested in studying two cases of free services and invalid/redirected application links. Specifically, we consider the following two instances: (1) Given certain type of free data services (e.g., DNS service discovered by our study) offered by operators, is it possible to exploit it to evade charges for other data services? (2) What is the current charging policy for application-level signaling or commands, which do not contribute to real content? Cases include FTP signaling over port 20, invalid HTTP links, HTTP redirects, and Email/IM signaling, etc. Note that in the second instance, these signaling messages are not the actual content. Whether to charge it or not depends on what view operators hold. Operators have every reason to charge it or not to charge it; it is not a right/wrong issue to address. But as the charging policy may also evolve towards more content based; this is an interesting topic for future study.

The “*how*” issue concerns the accounting architecture. It is about *how* to record mobile data usage in reality. As described in Chapter 2.2.3, the current 3G architecture takes the SGSN/GGSN based accounting approach. It records how much data volume has traversed the intermediate SGSN/GGSN inside the 3G infrastructure (see Figures 2.1 and 2.3 for an illustration). This element-based charging takes the local view inside




Figure 4.1: End-to-end data delivery path and involving components.

the 3G infrastructure, without coordinating with end devices when making accounting decisions. On the other hand, the data delivery path is always end to end. The end systems (e.g., the UE device or the server) may record usage volume different from what is logged inside the network. Note that the 3G accounting system does not explicitly collaborate with the end systems in its charging decisions. The effect is hence visible when failure or misbehavior occurs over the full delivery path.

In general, the end-to-end data delivery path consists of all six components (shown in Figure 4.1): the UE, the RAN, the 3G core network (CN), the middlebox, the wired Internet, and the host or server. Note that, the 3G network may deploy middle boxes (e.g., proxy servers, NAT boxes) over the delivery path, as shown in [WQX11], and SGSN/GGSN resides at the CN. Assume that SGSN/GGSN and CN are always functioning. Any other component may fail. Specifically, we consider four cases (1) The path segment between UE and RAN (i.e., the wireless delivery between BS and UE) experiences problems in delivery; (2) The path segment between the Internet and CN has packet drops; (3) The path segment between the middlebox and the host breaks; and (4) The host or server is not accessible.

The “*robustness*” perspective focuses on the proper/reasonable user behaviors. Also, cellular network infrastructure is not compromised. We are interested in the unexpected findings when normal (expected) actions run in possible scenarios.

The “*security*” perspective exploit the loopholes in the accounting system and launch charging-related attacks against the operator and (or) mobile users. Here, we

assume that the 3G charging subsystem is not compromised, i.e., all charging elements are operating properly (as described in Chapter 2.2.3). This implies that, the data usage records kept at SGSN/GGSN are not attacked, the mappings from CDRs to the flow, the bearer, and the mobile user are also intact. Moreover, user authentication within 3G/4G cellular networks works properly. Attackers cannot spoof other UE devices to access data services. In this category, we present two attacks: one is an undercharging (free, in fact) attack that exploits differential charging policy, and the other is an overcharging attack that injects any arbitrary volume of data to the victim. The latter loophole is rooted in the 3G/4G accounting architecture.

4.2 Experimental Methodology

We take an experimental approach to assess mobile data accounting in various scenarios. We design and conduct a series of experiments to examine the difference between the data volume recorded by operators and the one logged at the end device. In each experiment, we establish end-to-end data sessions from mobile phones to popular Internet services or our deployed server. We then record the data volume charged by operators and the ones observed at mobile phones or servers. We run main tests with two major mobile operators in the US, denoted as Operator-I and Operator-II for privacy concerns. They together offer nationwide coverage for 102.3M users, thus claiming about 50% of US market. For verification purpose, we also run similar tests with the third major US carrier that claims to support 4G LTE, a major carrier in China and Taiwan each.

Our mobile devices use several Android phone models: HTC Desire, Samsung Galaxy S1/S2, and Samsung Stratosphere (that supports 4G LTE), running on Android 2.2, 2.3.4, 2.3.6 and 2.3.5, respectively. Our experiments show that all the findings are phone platform independent; this is not hard to understand. We use an ASUS EeeBox PC EB1501 desktop as the deployed host outside cellular networks. It runs on an Intel Atom N330 1.6 GHz Dual Core processor and 1.5 GB DDR2 memory. It acts as a

content server (e.g., Apache web server, a FTP server using WingFTP software [Win], and a TCP/UDP server), proxy or attacker in various tests.

We use two methods to obtain data usage logged by operators. The first one is to dial a special number to retrieve the remaining monthly data usage in a near real-time mode. Most operators support this Dial-In feature, e.g., via dialing #DATA for Verizon, *DATA# for AT&T, and #932# for T-Mobile in the US, sending a a short message of CXLL for China Mobile. The data usage will then be delivered via a text message after this Dial-In. By logging data usage volume before and after our experiment, we obtain the usage volume observed by the operator during the experiment. The second method is to log onto the mobile operator website and access online data usage records (as shown in Table 4.1). Operator-I only supports DIAL-IN method, while Operator-II supports both. We thus use the first method for Operator-I and the second method for Operator-II. In terms of report latency, we find that Operator-I may report data usage in five to ten minutes while Operator-II may take up to six hours to update their records. However, Operator-II provides an itemized data charging volume associated with the timestamp; a new item will be generated when a new PDP context (bearer) is established. We thus conduct experiments with proper time window (> 10 minutes) and establish a new PDP context for each experiment to avoid confusions and cope with latency. Both operators support 1 KB accuracy in data usage report. Since data usage logged by the operators only has timestamps, we need extra mechanisms to ensure that data usage belongs to the specific application or data flows in our experiments. For this purpose, we clean up the runtime environment (factory reset and disable "Background data" and "Auto-sync" functions). We also use monitoring tools (Wireshark [Wir]) to capture all-level packets to/from the phone. Whenever unwanted services observed, we rerun the experiment. For those tests that last for an extended period of time (e.g., three or six hours), we use Wireshark to filter out those unrelated packets in the analysis.

We use two tools to log data usage on mobile phones. The first is to use Traffic-Monitor [Traa], a software tool available from Google Play. It records data volume for

each application with 0.01KB accuracy. The second is to use our own tool, which is written via the TrafficStats class interfaces [Trab] in Android SDK to collect network traffic statistics. We record the number of packets and bytes transmitted and received on all interfaces and on a per-application basis. We use both tools to record the UE data usage and verify whether the usage is consistent or not. We further use Wireshark to log the traffic statistics at our server.

We run experiments in cases of normal settings and extreme scenarios. The normal settings capture users' common usage patterns, including popular protocols such as TCP, typical applications, and daily-life usage. The extreme scenarios are carefully created in experiments, and seek to stretch out the accounting system in worst-case settings. The extreme scenarios also cover the cases where the accounting attacks are launched. In each experiment, we record the data usage from operator V_{OP} , the one observed from mobile phone V_{UE} , and the one at server V_{SR} if used. We conduct each experiment for 5 –15 runs. The experimental results are quite stable in different runs. Unless explicitly stated, we show the average values, not individual ones.

In the following chapter, we present our findings in various experiments. We will see that the usage-based accounting scheme has been shaped by both the 3G standards and the operators' policy practice. The standards define the architecture and mechanisms, whereas the operators specify their own charging/accounting policies. The former determines *how* to do accounting and the latter define *what* to be accounted.

4.3 We Get What We Want For Free

The first finding is that we can get what we want for free. The root cause is that, the current charging policy practiced by operators has loopholes, and can be exploited to build "free" data services. Our study shows that, both operators offer free Domain Name System (DNS) service via transport-layer port number 53. There is almost no enforcement mechanism to ensure that the packets going through this port are indeed DNS

messages. Even worse, no effective mechanism exists to limit the traffic volume going through this port. Consequently, this free service can be readily abused to create “toll-free” data services. We have used three approaches to build a simple “Toll-Free-Data Attack” prototype, and demonstrated that it is feasible to offer various data services, e.g., file downloading or video streaming, through a special proxy server relaying data over the free transport-layer port. The process is similar to calling 800-like voice hot-lines, but for free data access. Finally, we make suggestions to fix this “bug.” Note that, operators have fixed these loopholes once we published the results; this attack is not valid any longer.

4.3.1 Loopholes in Charging Policy Practice

The 3G standards offer the operators flexibility to define their own policies on what to charge. Unfortunately, their policies and implementations may contain serious flaws.

We use the example of Web browsing (`www.cnn.com`) to illustrate the vulnerabilities in the charging policy practice. Figure 4.2 illustrates typical steps for Web browsing. Upon receiving the target URL, the Web browser immediately initiates two actions. One is to send a DNS query to request the IP address for this URL. The other is to send a HTTP query to the Web server using the obtained IP address and receive a HTTP response. In mobile data charging, the above operations invoke two charging flows. One is the DNS query/response which goes through the CN to the DNS resolver or server. It is primarily carried by UDP on port 53, though TCP over port 53 is also allowed [RFC10]. The other flow is for HTTP, which traverses the CN to enable communication between the UE and the Web server. It runs on TCP using port 80 (or other ports, e.g., 8080 or 443 for https). The CN records the data volume associated with each flow for billing.

Our study shows that, both operators tested in our experiments offer free DNS service. This policy practice makes sense, since DNS is considered a fundamental service




Figure 4.2: Procedure of a typical mobile Web browsing.

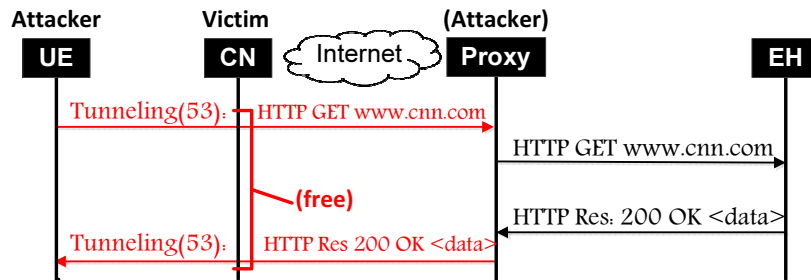


Figure 4.3: Illustration of free Web browsing in the “toll-free-data-attack”.

to jump start Internet applications. Almost no Internet services can be initiated without DNS. DNS service is offered for free by many public DNS servers (e.g., Google, OpenDNS [free]). Operators thus have every reason to offer it for free, to facilitate followup data usage by other Internet services. Therefore, free DNS service can be well justified as a good (at least reasonable) policy.

However, the operator practice to offer free DNS service does have two loopholes. First, there is almost no enforcement mechanism to ensure that the packets going through this DNS-reserved port are indeed DNS messages (*free fake DNS loophole*). Second, there is no effective mechanism to limit the traffic volume going through this port (*no volume-check loophole*). We next elaborate both loopholes using experiments.

Free fake DNS loophole Internet RFC 5966 stipulates that DNS service is offered using transport-layer port 53 via UDP or TCP [RFC10]. To identify a data flow, the 3GPP standards define five-tuple flow ID composed of source and destination IP addresses, source and destination port numbers, and protocol ID (see Chapter 2.2.3).

However, both operators do not strictly enforce this service via the standard five-tuple flow ID, but via only the destination port (plus protocol ID for Operator-II), , thus exposing an obvious vulnerability.

We use experiments to verify whether the DNS service is free and what exact factors the free DNS service depends on in the operator’s implementation. We conduct five experiments: (1) *DNS-Default*: send 100 DNS queries to the default DNS server provided by the operators; (2) *DNS-Google*: send 100 DNS queries to a Google DNS server (IP address: 8.8.8.8); (3) *TCP53-Google*: repeat (2) but via TCP at port 53; (4) *TCP53-Server*: send 50 random packets to our own server using TCP via port 53, and request the server to return the received packets; each packet (including IP/TCP headers) is 1KB; Source port number is randomly chosen; and (5) *UDP53-Server*: repeat (4) but using UDP. The goal of these experiments is to verify what factors the free DNS service depends on: (a) Does the free DNS service depend on the destination address? For example, is DNS only free via the operator DNS resolvers/servers, see (1) and (2)? (b) Does it depend on the protocol ID? For example, is it free for both UDP and TCP, see (2) and (3). (c) Does it depend on the source port number or check the DNS message semantics, see (4) and (5)?

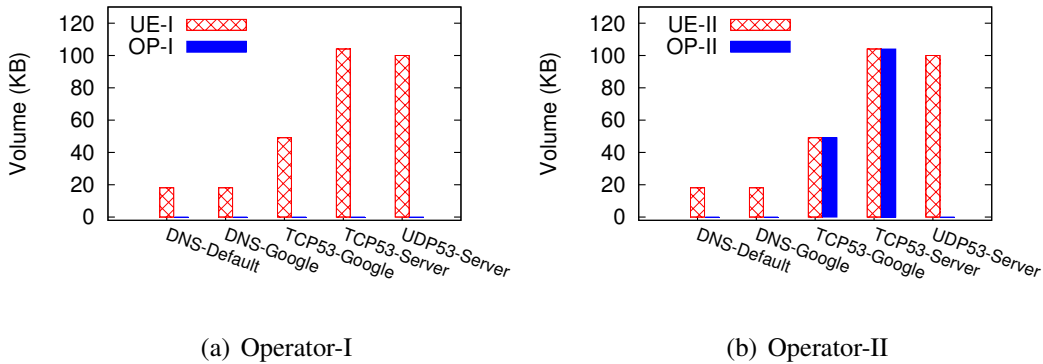


Figure 4.4: V_{UE} and V_{OP} in five DNS tests.

We conduct these experiments with two US major operators. We have purchased unlimited daily data plans from both operators and thus do not run into legal issues while testing free data services (the actual data usage is not counted by operators). We

invalidate the hypothesis that the operator has no incentive to correctly report the traffic usage for users with unlimited access. To this end, we also use 200MB and 4GB data plans in free data service tests, and compare the results with using unlimited data plan. Results are consistent in all three plans. We further test different services (e.g., Web, YouTube, Gmail) using our unlimited data plans and verify that the data usage records at the UE and the operator are consistent. Figure 4.4 plots the data volume observed by the UE and two operators in all five cases. The results show that,

Operator-I: Packets via port 53 are FREE
Operator-II: Packets via UDP + port 53 are FREE

Specifically, the UE sends and receives about 18.1 KB for 100 DNS queries and responses in both DNS-Default and DNS-Google tests. In the TCP53-Google test, the traffic volume rises to 48.1 KB due to TCP signaling overhead (SYNC, etc). In both TCP53-Server and UDP53-Server tests, the UE sends and receives 100 KB as expected. Operator-I charges for free (i.e., $V_{OP} = 0$) in all cases while Operator-II charges those TCP cases. From these results, we learn that the free DNS service is implemented by Operator-I using only one field in the flow ID (i.e., the destination port 53). In contrast, Operator-II enforces free DNS service using two tuples in the flow ID, i.e., UDP over destination port 53.

No volume-check loophole Our study further shows that, there is no mechanism to limit the traffic volume going through this free-service port. To this end, we build our own server outside the cellular network that exchanges data services with mobile phones using UDP over port 53. For operator-I, we also enable toll-free TCP at port 53. We perform three experiments: (I) *Free-One*: the UE sends one request to our server to download a 5MB file; (II) *Free-Equal*: the UE uploads a 3MB file to our server, and requests to return the delivered packets; (III) *Free-Long*: the UE sends many small requests (100 B) to our server for an hour, each of which requests a 1KB response. These experiments are to validate whether the free data service can support unbounded traffic upon a single request, whether it supports large-volume uplink and downlink

traffic, and whether it allows for long-lived data access.




Figure 4.5: Feasibility test of free data services; $V_{OP} = 0$.


Figure 4.5 plots the data volume observed by the UE and both operators in the above three scenarios. It shows that, both operators can be exploited for free data services in all these scenarios, except that Operator-I does not allow unbounded traffic for one fake “DNS” request. The fake DNS message In the first test, Operator-I only allows to deliver 29 KB downlink data to the UE, while Operator-II delivers much larger file (up to 4 MB). We gauge that Operator-I might have enforced a checking mechanism to verify the size of the response message, in which a real DNS message size is typically bounded. However, this size checking can be easily bypassed. The UE simply sends out many small, dumb packets over this session, to increase the quota for downlink traffic. Then large downlink data can pass this checking. This has been validated in scenarios (II) and (III). In these tests, the gap between V_{UE} and the expected file size is mainly caused by unreliable transmission via UDP. These results demonstrate that free DNS service is readily abused to create any “free” data services. Since each scenario has ten or more experimental runs at the data rate from 100 Kbps to 1 Mbps, the total free data we have obtained from this DNS hack exceeds 200MB. In fact, we believe that it can support any amount we specify.

4.3.2 Toll-Free Data Service Attack

We now show how to exploit above two loopholes to launch “free” mobile data access attack. The key idea is to use a proxy server (placed outside the cellular network) to bridge the data access between the mobile phone and the Internet server. The communication between the proxy and the phone is carried out over the free channel (i.e., UDP or TCP over port 53, depending on the operator policy). We use “tunneling” between the UE and the proxy server, encapsulating data packets in DNS messages, which traverse the 3G network free of charge. The proxy server relays packets on behalf of the UE. Free communication is thus extended to between the UE and an Internet host, while the 3G core network (CN) is the victim. Figure 4.3 illustrates the example of how Web browsing becomes free of charge. The process is similar to calling 800-voice hotlines, but for free data access. We thus name it as the “*toll-free-data-access-attack*.”

We present three approaches when implementing the toll-free-data-attack. All show that, it is simple enough to obtain free mobile data access in reality. The first approach is to use a HTTP proxy running on port 53. It is easily done using available free proxy software such as FreeProxy [Fref]. The mobile Web browser is then configured to use the established HTTP proxy, as shown in Figure 4.6(a). This approach is easy to implement; no coding and hacking are needed. However, it only works for Web browsing and for Operator-I, which allows free TCP via port 53. To evaluate its effectiveness, we test two Web browsers – Mozilla Firefox and Opera Mobile [Mobb], one hour each. We are able to use Operator-I network for free, while the actual data volume goes beyond 20 MB.

The second approach is to use a socks proxy. It works with various application protocols, e.g., HTTP, FTP, SMTP, POP3, NNTP, etc. Similarly, we deploy a socks proxy running on port 53. On the phone side, we install ProxyDroid [Pro] to enable socks proxy functionality. The phone configuration is shown in Figure 4.6(b). This method supports more applications without configuring each application individually.



(a) HTTP proxy

(b) Socks proxy

(c) FDP proxy

Figure 4.6: Three approaches to launch “toll-free-data-access-attack.”

However, it still only applies to the TCP-53-free operators. We assess this attack with Operator-I using mobile applications, e.g., Web browsing, YouTube, Gmail, Google Map, Skype and FTP (via AndFTP [And]). The results show that, all services are free of charge except Skype voice call and FTP download. We figure out that, these two applications fail to go through the socks proxy; It is an implementation issue in ProxyDroid.

The third approach is to deploy a proxy server to enable “tunneling” between the phone and itself. To this end, we design a Free Data Protocol (FDP) to encapsulate data packets between the UE and the proxy into fake DNS messages, i.e., to carry packets in ANY-on-port-53 flows for Operator-I and UDP-on-port-53 flows for Operator-II. These messages are any data packets, not following DNS semantics. To bypass the limit of data volume for one fake DNS request (for Operator-I), FDP also periodically sends small KEEP-ALIVE messages from the UE side. The attacker enables the FDP at the UE and the proxy server. Note that, the DNS-tunneling idea is also used in the iodine [Iod] and NSTX [NST] tools to enable Internet access over DNS. Moreover, the NSTX was used to demonstrates the similar idea for free Internet access with a toll-free Microsoft PPP dial-in number in Germany [ip]. Both work in the wired Internet and free Internet access is available with specific DNS servers. In our experiments, we have built a simple prototype that revises applications to use FDP. We test our prototype with the revised HTTP and FTP applications working on top of FDP. Figure 4.6(c) captures

the screen shot when visiting `www.cnn.com`. It shows that, data access is free for both operators while the actual data volume reaches 100 MB. Moreover, the upper limit of free traffic volume seems unbounded in our tests.

4.3.3 Policy as Double-Edged Sword

Policy practice is an inherent component of the accounting systems for mobile users. Policies can be good for both operators and end users. On one hand, policy practice offers carriers flexibility, while injecting dynamics into the market. It can serve as a viable mechanism to compete with other carriers when offering users better services at lower cost during certain times. On the other hand, users can also benefit from policy practice. As we have seen in the DNS case, users will pay less due to the free DNS service!

In addition to DNS hacking, other tricks exist for free data services by exploiting the loopholes in the charging policy. For example, in case some operators offer free Internet access to a given website (including free access to mobile Facebook [frea] or a given Web site [fred]), Web redirection from one free Web server to the target Web page is used to enable free data services; using certain, free Access Point Name (APN), which is a configurable network ID used by a mobile device when connecting to a carrier, offers another way for free data service in some carriers, e.g., AirTel India [frec] and UK Three [free]. These examples, including our DNS hacking, show that *differential-charging* policy offers flexibility but may also be abused if not enforced properly.

Fundamentally, for a metered charging service, people necessarily have incentives to exploit and abuse any transfer that is free or cheaper. However, if history is the best teacher (BGP is a good example where people learned painful lessons with its policy-based routing), we have to be prudent with policy practice. The policy choice needs to be conflict free. Moreover, its enforcement has to be strict. Otherwise, policy may open holes that operators may never anticipate.

4.3.4 Recommended Quick Fix

The simplest fix is to stop offering free DNS service or other forms of free data services that can go outside cellular networks. People always have incentive to abuse any differential-charging services; Therefore, the simplest, possibly also the best solution to abuse prevention is to eliminate the free services. Moreover, DNS traffic is negligible in normal cases; it should lead to no noticeable difference in most usage scenarios.

We also seek remedies to fix this bug while still retaining the free DNS service. For example, we have considered that the operator can provide quota for free DNS service. The DNS data usage beyond the quota will be still charged. Ideally, the quota should be assigned based on the average usage patterns. It can be a fixed amount or a percentage of the data usage. The challenge for this approach is how to set an appropriate quota. Some applications or services such as MobileMe [moba] and DNSSEC [dns] may heavily use DNS while others do not.

The alternative approach is to enforce checking on the destination IP address of the DNS request. For example, free DNS services are only allowed when these messages go to designated or authenticated DNS resolvers or servers managed by carriers. This enforcement eliminates the possibility to go through those fake DNS servers. The possible downsides include: extra effort is needed to authenticate DNS servers, not all DNS servers across the Internet can be directly accessed by UE, and workload at the designated DNS servers increases. Unfortunately, it is still possible for malicious users to deceive those resolvers/servers to forward fake DNS requests to a fake DNS server. The only difference is that the attack cost could be higher.

In the more general context, when the charging policy allows different unit-prices for diverse services, extra bullet-proof mechanisms are required; otherwise, the attacker always seeks to use the cheaper one. However, the deployment and operation of such security mechanisms will inevitably increase the cost of the carrier. Moreover, the security mechanism still needs to ensure itself to be secure in its design and operation.

All these pose interesting research issues for the future.

4.3.5 Progress Update and Carriers in Other Regions

We run similar tests with other carriers. We indeed observe that the free DNS policy be operator dependent. The third US operator also offers UDP-based DNS for free, and behaves similar to Operator-II. However, for both carriers in China and Taiwan, the DNS service is not free. Both operators charge DNS messages identical to data traffic.

When we disclosed this finding in 2012 [PTL12, PLT12] (probably July/August), all these three US operators took actions and stopped free DNS services. As a result, such toll-free data access is not valid now. This finding is only used to demonstrate the loopholes in charging/accounting policy and remind us to be cautious of making a policy that is designed with good intention but may not work as expected in reality.

4.4 We Pay For What We Do Not Get

The second finding is just the opposite. We might be charged for data that never reach us or the data we never deliver to the destination. The root cause is that, the data volume is recorded inside the cellular network core without taking feedback from the end device when making accounting decisions; it can be different from the volume received at end device. We first describe the results in the extreme cases, which represent some worst-case behaviors and may rarely occur in reality. We then describe the average cases, which show how applications and users behave in common usage scenarios. We elaborate on how large the difference between the user's usage and the operator's charge can reach in these scenarios, explain their root causes, and suggest quick fixes.

4.4.1 Extreme Cases

We first examine how bad overcharging can become in certain extreme conditions. The goal is to expose the potential downside of the largely successful 3G charging system. Note that these conditions do not represent the typical usage patterns in practice. They could occur in reality but only infrequently.

In these tests, data traffic is delivered using UDP between an Internet server and a mobile phone. Though most network applications run on TCP, recent traffic study [Cai] shows that, UDP is still used for data delivery in 10-20% applications, including video streaming, VoIP, and Virtual Private Network (VPN), etc. We consider downlink cases in both no-signal zone and weak-signal zone for the mobile device; these zones vary with locations. Our experiments show that, they are mainly caused by poor coverage by carriers. For example, we have experienced three to four dead zones or zones with very weak signals measured by RSSI on our office floor of the campus building. Mobile devices are unable to receive data when suddenly entering into a dead zone without signals. However, they are still charged though such data never reach them. In the worst-case scenario, we have observed that charging proceed for more than three hours and result in more than 450MB data if the application has no control loop!

4.4.1.1 UDP in No-Signal Scenario

We conduct DL-NS experiments to put our phone into a dead zone without signals, and see what happens to the ongoing downlink UDP transmission from an Internet server to the UE. The goal is to examine whether the data usage charged by operators differs from that received by mobile phones.

Our experiments are conducted in an indoor environment shown in Figure 4.7(a). The coverage varies at locations for both operators. Figure 4.7(b) plots the medium of the measured received signal strength indicators (RSSIs). RSSI values vary from -

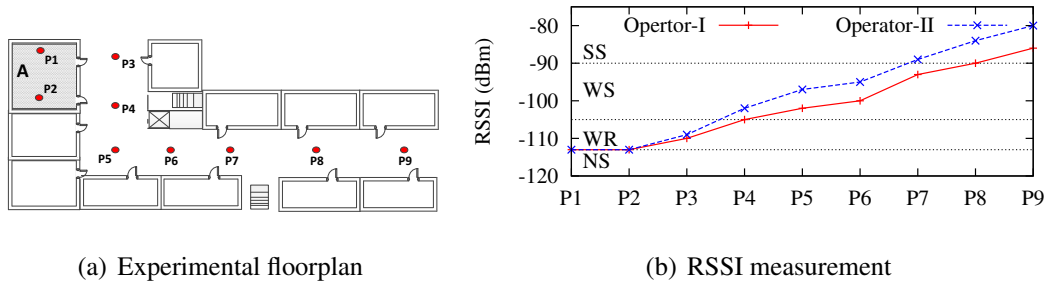


Figure 4.7: Our indoor testbed, where room A is a dead zone (NS-zone) for both operators.

113 dBm to -80dBm ¹ at various spots and fluctuate within 3 dBm at each spot. Based on RSSI values, we divide the whole area into four zones: (1) *SS-zone* with strong signals (RSSI > -90 dBm); (2) *W-zone* with weak signals (-90 dBm \geq RSSI > -105 dBm); (3) *WR-zone* with weaker signals (-105 dBm \geq RSSI > -113 dBm); and (4) *NS-zone* (i.e., dead zone) with no signals (RSSI \leq -113 dBm). Note that, different operators yield different coverage strength; Operator-II has stronger signal strength than Operator-I in this setting. However, Room A remains a NS-zone for both operators. We also conducted prior experiments (e.g., making a phone call) to ensure that the phone is indeed out of service in Room A.

We now describe how to set up the DL-NS experiment step by step, as shown in Figure 4.8. First, at P9 (i.e., in the SS-zone), we send a UDP request from the mobile phone to our own server to start this experiment; Once the communication is ongoing, the server responds with an acknowledge message to the phone and sets a timer, which triggers UDP data transmission upon timeout. Upon receiving the ACK, we move the phone from the SS-zone to the NS-zone (i.e., Room A) (Step 2), hopefully before timeout. Since it takes about 30 seconds to walk into Room A, the timer is set as one minute to keep the server stay idle (no data delivery) during Step 2. Upon timeout, the server transmits UDP packets to the phone at a constant data speed s for another t

¹ -113 dBm is the lowest signal strength that a typical mobile phone can receive; it implies that the phone is out of service. In indoor environment, strong signal strength is much smaller (here, $(-90, -80)$ dBm) than the outdoor one that usually reaches -65 dBm. dBm stands for the measured power ratio in decibels (dB) referenced to one mW and 0 dBm equals 1 mW.

minutes (Step 3); s and t are configurable parameters in the experiment. During Step 3, the phone remains in the NS-zone. We record data usage V_{SR} , V_{UE} , and V_{OP} , observed at the server, the UE, and the Operator, respectively.




Figure 4.8: Procedure of DL-NS experiment.

Results We first set the UDP source rate as $s = 50$ Kbps and the data transmission lasts for $t = 10$ minutes. Our server sends about 3.75 MB data ($50K \times 10 \times 60/8 = 3.75M$), similar to the volume charged by the operator (3.73 MB). The minor difference between these two volumes (i.e., $V_{SR} - V_{OP}$) is mainly caused by occasional packet loss. However, the mobile phone does not receive any such data, except the 80 B for one UDP request and one ACK message at the start. This result shows that the charging infrastructure could charge mobile users of data that never reach them in case of a UDP-based application without control loop. Moreover, we believe that, many mobile users might not be even aware of such a charge. It is quite common that mobile phone users unconsciously enter into a NS-zone in reality. They have no clue that roaming into the no-signal region may incur data volume charge by the operator, if the UDP sender is still transmitting.

4.4.1.2 Worst-Case Observations

We test with various source rates and different durations. The gap between the operator charge and the volume received by users (i.e., $V_{OP} - V_{UE}$) can be approximated by $s \times t$, which is exactly the volume of data sent by the server but never reached the phone:

$$\text{Accounting-Volume-Gap} = V_{OP} - V_{UE} \approx s \times t.$$

Our experiments show that the approximation still holds even when the speed s goes up at least 8 Mbps or the duration t lasts three hours²! The above finding shows that, the operators charge mobile users based on the data volume sent by the server and arriving at the cellular core network, but not the volume that cellular networks have actually delivered to the users. This rule still applies no matter how large the gap could turn into. For example, the operators have charged us for 450 MB in one run when the server keeps sending downlink data at 1 Mbps for one hour ($1 \times 60 \times 60/8 = 450M$), even though no single data bit arrives at the mobile phone!

We change the source rate s from 50 Kbps to 8 Mbps to examine how the gap varies with high data speed. Figure 4.9 plots the Accounting-Volume-Gap (called “gap” hereafter) for Operator-I using two sending servers with different link capacities. The transmission lasts one minute. The results are similar for Operator-II³. Note that in DL-NS experiments, the UE receives almost zero bits and the gap is approximately equal to the volume charged by operators V_{OP} . It is seen that the gap is in proportion to the UDP source rate s in Server-1 case (in Figure 4.9(a)). For Server-2, we find out that the gap is almost the same as V_{SR} when the data rate s is low (≤ 2 Mbps); when the source rate increases (> 2 Mbps), the one charged by operator is smaller. This is because Server-2 uses home Internet service and has bounded uplink speed. In contrast, using Server-1 with higher uplink bandwidth, the operator charges us for about 58.7 MB (close to $V_{SR} = 60\text{MB}$) in one minute at the 8 Mbps rate. This test infers that, the operator

²It holds if all the data can arrive at the operator without packets loss in the Internet.

³The results for Operator-II will be omitted hereafter to save space if they are similar.

charging practice is only based on how much data would arrive at the core network, no matter how fast it is. Without much packet loss or congestion, gap grows in proportion to the UDP source rate.

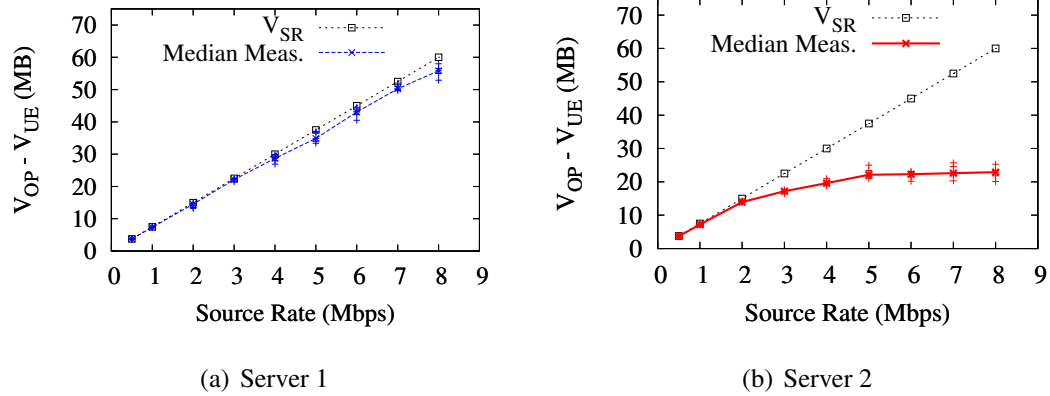


Figure 4.9: Accounting volume gaps under various UDP source rates; $t = 1$ minute.




Figure 4.10: Accounting volume gaps with various in-NS-zone durations.

Even worse, the operator may charge us for a long time. We put the phone in the NS-zone for different durations to see how long the gap may last. If the application layer tears down the session once the phone is out of service for a small duration of time, the gap would be small and not incur a large bill. However, if the application does not terminate, we find out that gap could last at least **three hours!** We run experiments for the slow session (50 Kbps) up to six hours and the fast session (1 Mbps) up to one hour. Figure 4.10 plots the gap when the in-NS-zone duration varies from one minute to

six hours. Within the initial three hours, the gap for the slow session grows linearly with the duration t . Both operators stop recording when the usage reaches about 66.3 MB, which approximates about three-hour data transmissions. We do not run experiments for high-speed UDP sessions (e.g., 1 Mbps or even 8 Mbps) up to three hours, because the data usage probably goes up to 1.35 GB or 10.8 GB, which incurs a huge bill. In fact, the gap as large as 450 MB and the charging duration of about three hours are already significant enough.

4.4.1.3 Still-Bad Case: Even With Signals

We next show that, the charging gap still exists even when the wireless link is not broken. The gap concerns the wireless environment in terms of available radio link rate. We conduct another DL-ALL experiment, where the mobile phone is statically placed in different zones with various signal strengths. Different from the DL-NS experiments, UDP packets are immediately transmitted once the handshake between the phone and the server is established.

Figure 4.12 plots the gap when the phone is placed in zones with different signal strengths under various source rates in Operator-I. Each data transmission lasts one minute. The figure shows that, the gap becomes larger as the signal strength becomes weaker or as the source rate becomes larger. Figure 4.11 illustrates why it happens, and Table 4.3 shows the detailed results for three examples of experimental traces.




Figure 4.11: Illustration of gap creation in various wireless environments.

We make three observations. First, the core network receives almost all data packets (i.e., $V_{OP} \approx V_{SR}$). The charging gap is still caused by the unsuccessful packet delivery




Figure 4.12: Gaps under various source rates in different zones in DL-All experiments.

Setting	V_{UE} (MB)	V_{OP} (MB)	V_{SR} (MB)	Link rate r (Kbps)	Finish time (sec)	gap (MB)
SS-zone (-84 dBm)	6.0	6.0	6.0	644.4	74.7	0
W-zone (-98 dBm)	2.90	6.0	6.0	326.7	71.0	3.10
WR-zone(-109 dBm)	1.46	6.0	6.0	168.1	69.5	4.54

Table 4.3: Example results for three DL-All experiments when source rate is $s = 800$ Kbps and $t = 1$ minute.

from RAN to UE. Second, packets are dropped in RAN because the incoming source rate is much higher than the effective rate r of the wireless link to the phone (see Figure 4.11). The effective rate depends on wireless signal strength. For example, the effective rate is 168.1 Kbps in the WR-zone, much smaller than in the SS-zone (about 644.4 Kbps). Third, not all the mismatches between the source rate s and the effective rate r lead to packet drops. Take the example of SS-zone with $s = 800$ Kbps. It spends more time (about 74.7 seconds) and incurs large delay. We infer that this attributes to the buffer mechanism, which temporarily stores incoming packets (if too fast) and retransmits them if needed. However, as the source rate further increases, the speed mismatch becomes too large (especially, in the WR-zone/W-zone) to be handled by buffers, leading to eventual packet drops. Consequently, we still pay for bits that never reach us even though wireless links exist. The charging gap depends on the operating

environment.

We also consider the case with intermittent signals where mobile users may lose signals for a while but recover them shortly. This scenario is common with cases of mobility and special landscape (mountains or high buildings). Our findings show that, those packets that the phone miss in NS-zones still contribute to the charging gap, though the communication recovers soon and buffering and retransmission mechanisms reduce the charging gap to some extent. Figure plots the gap when the phone loses signals for 10, 30, 60, 90 seconds. In the meantime, the UDP server sends packets at a speed s . We can approximate the data volume that arrived in t time but finally received by the phone is $V_{back} = s \times t - (V_{OP} - V_{UE})$. Figure 4.13 plots $(V_{OP} - V_{UE})/s$ (i.e., $t - V_{back}/s$) under various in-NS-zone durations. The larger the duration, the fewer the received packets. The results imply that, buffering and retransmission do contribute to packet delivery (retrieving 15 out of 90 seconds data in 50 Kbps-UDP session). However, those packets not recovered are still counted into the mobile bills. It also shows that the gap exists even when mobile users only lose signals for several seconds.

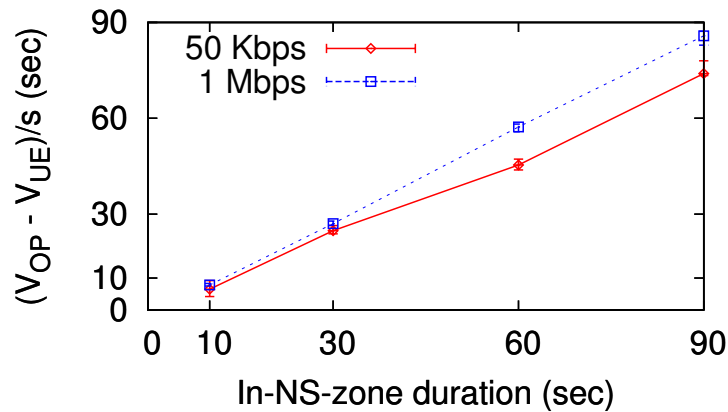


Figure 4.13: Gaps with with intermittent connectivity in DL-All experiments.

4.4.2 Root Cause in Open-Loop Data Accounting

We now explain the root cause. It lies in the open-loop data accounting architecture.

In the downlink case, traffic is delivered from the external server to the mobile device via cellular networks (e.g., GGSN, SGSN and RNC in turn). It is easy to see that, the observed data volume monotonically decreases along the downstream delivery path, i.e.,

$$V_{UE} \leq V_{SGSN} \leq V_{GGSN} \leq V_{SR}. \quad (4.1)$$

Due to unreliable packet delivery, packets might be dropped at any intermediate node, thus incurring the volume gap. For example, in DL-NS experiments, the last hop is broken, so no data would be delivered to the phone ($V_{UE} \approx 0$). However, the 3G/4G accounting system obtains data usage based on the volume recorded by SGSN and GGSN. Therefore, those UDP packets, which arrive at GGSN or SGSN but never reach the UE, are still counted as the data usage by this UE. This results in a large gap between the actual data usage and the billing volume.

Note that the open-loop accounting works well for voice calls. Cellular networks also use a usage-based charging for voice calls (e.g., how many minutes are used). But the difference is, voice calls use circuit-switched technology. With rich signaling and control, network equipments have a consistent view of call status. For example, when the phone is in a no-signal zone, the voice call will drop and the drop status will be synchronized among all the network parties; Thus, the core network will sequentially stop the accounting. However, for data transmission, it is not true. Data delivery is based on packet-switched technology where each hop is independent. As a result, along the end-to-end path, network equipments may have inconsistent views on connection status and usage, if without any coordination or feedback. The open-loop accounting architecture lacks coordination with end systems and thus the local decision made in the core will differ from the end view. Moreover, this implies that a successful design for voice may not be applicable to data. Our intention to inherit nice features in a telecom system may impose risks for the Internet-like design. The add-on design might become problematic for a large-scale, complicated system.

4.4.3 Common Cases

We now study the common cases, which reflect the usage patterns by applications and users in their daily activities. Our study has three categories. The first is to see how TCP, the dominant transport protocol for applications, reacts in the no-signal and weak-signal zones. In the second category, we study five popular applications, including Web browsing, Skype for VoIP, YouTube, PPS streaming, and streaming over VPN tunnels. In the third category, we report the user-based, weekly accounting gap.

4.4.3.1 TCP Cases

We next study the charging behaviors of TCP flows, which turn out to be not too bad in terms of the overcharged volume. Intuitively, TCP behaves differently due to its built-in mechanisms of congestion control and reliable data transfer. Its feedback loop offers implicit coordination between the network and end devices. We conduct DL-NS experiments via TCP, but let the server deliver packets without timeout. We keep the mobile phone longer in the SS-zone before roaming into the NS-zone. As expected, the gap greatly reduces; it is seen to vary between 2.9 KB and 50 KB in our experiments. As we know, the charging gap is determined by how many bytes are delivered before automatic TCP session teardown. This is determined by the congestion window size before the UE enters into the NS-zone and the timer for automatic teardown.

When the phone is out of service, packets in the congestion control window are still sent out; Since no more packets can be acknowledged and the window decreases; the unacknowledged packets are retransmitted until automatic connection teardown. Note that, some ACKs at the UE fail to be sent out due to the broken connection. Thus, the gap is calculated as the actual data volume sent by the server but not received by the UE, i.e., $V_{InFlight} + V_{ReTx} - V_{ACKUnsent}$. This is typically bounded to tens or hundreds of KBs by mobile TCP configuration. Figure 4.14 plots an example of TCP segment records observed by the two ends (the gap is 45.7 KB here). Note that the ACKs for




Figure 4.14: One DL-NS experiment trace using TCP.

those packets received by UE may not be sent out due to the sudden connection break, the actual data packets sent by the server but not received by the UE is calculated by $V_{InFlight} + V_{ReTx} - V_{ACKUnsent}$, which equals the gap.

In reality, most current Internet applications are built on the top of TCP. The inherent congestion control and automatic connection reset may release this connection once it fails for an extended period of time. This is why the charging gap is not large in reality. The flaw in open-loop data accounting is mitigated by external TCP control. However, this is not the best practice from the 3G accounting perspective. It has to rely on higher-layer protocols such as TCP to handle abnormal behaviors of both connection teardown and overcharging. Moreover, UDP-based multimedia or other applications may still suffer if they do not implement control logic for automatic session teardown. Since transport or higher layers are beyond the control of cellular infrastructure, they are vulnerable to any malicious revision that creates unexpected accounting gaps. The security problem will be addressed in Chapter 4.5.

4.4.3.2 Application Behaviors

We carry out DL-NS experiments using five applications, including Web browsing, Skype, YouTube, PPS streaming [PPS], and VLC streaming [VLC] over VPN tunnels.

In the web browsing test, we visit `www.cnn.com` at different times of the day; In Skype test, we make Skype video call, which uses UDP-based data delivery with built-in rate control [ZXH12]. Other three applications provide video streaming on phones. YouTube is TCP based, while PPS and VLC streaming are UDP based. PPS streaming is a very popular peer-to-peer application in China, whereas VLC + VPN offers one method to watch home HD TV. These applications have built-in control mechanisms, which can reduce rate or even tear down data delivery if the network connection degrades or breaks. We run each application for 5-15 runs with two US operators, except that VLC+VPN is blocked by the firewall of Operator-II. We start these applications in a SS-zone and enter into a NS-zone in several minutes.

Table 4.4 records the observed volume gap ($V_{OP} - V_{UE}$). The gap is negligible for Web browsing, and no more than 1MB for each Skype or Youtube run. However, it may reach up to 4.3 MB and 29.9 MB for PPS and VLC+VPN streaming, respectively. Interestingly, we observe that the gap be negative (but close to zero) for Web browsing; we figure out that it is because some packets (i.e., DNS packets) are free of charge, to be elaborated in Chapter 4.3. The difference between Web volume-gap of two operators is caused by different versions of Web pages. Mobile CNN page (about 0.2 MB) is fetched for Operator-II while the official CNN page (about 1.2MB) is fetched for Operator-I; the percentage of the volume gap is about 1.5–2.5%.

APPS	Operator-I			Operator-II		
	Med (MB)	Max (MB)	Min (MB)	Med (MB)	Max (MB)	Min (MB)
Web	-0.03	-0.04	0.00	-3KB	-4.6KB	0.6KB
Skype	0.88	0.99	0.40	0.68	0.99	0.10
YouTube	0.23	0.34	0.20	0.44	0.63	0.36
PPS	3.30	4.3	0.72	1.4	1.6	0.92
VLC + VPN	2.97	29.9	1.45	-	-	-

Table 4.4: Volume gaps for applications in DL-NS experiments.

We also find that the volume gap varies significantly even for the same application (except Web browsing). For example, the gap varies from 0.1MB to 0.99MB for dif-

ferent Skype runs in Operator-II. It turns out that the average transmission rate during the ten seconds before entering into a NS-zone varies from 221 Kbps to 1.0 Mbps. We observe that the gap be bigger if the transmission rate before going into the NS-zone is larger. It is easy to understand, since the volume gap is contributed by the source rate s and durations t , where s is the average source rate during the period in the NS-signal zone and determined by rate control for each application. We also observe large gap for VLC+VPN. This is because its automatic teardown timer is larger. The teardown timer in VLC+VPN lasts from 30 seconds to several minutes (about 6-minute value was observed in our experiments, leading to 29.9 MB charging gap), whereas it is merely several seconds in Skype. These application tests again demonstrate that, though mobile data charging is largely successful in practice, non-negligible overcharging is still observed due to problems in the current architecture.

We conduct another experiment to assess the performance over intermittent wireless channels. We watch videos via VLC streaming when we roam around the office area with several NS-zones. The wireless signals are intermittent, but usually recover within minutes. The video halts when we lose signals, but resumes once the wireless link is reestablished. In our experiment, we see that this video streaming never tear down and the observed volume gap reach up to 27.7 MB. Moreover, the gap depends on the number of SS-NS zone switches during our movement. We observe 11.8 MB and 27.7 MB charging gap for 10-minute and 30-minute movements, respectively. In DL-NS experiments, the observed gap for VLC streaming is at most 2.97 MB since the server tears down video streaming upon losing responses from the mobile client.

4.4.3.3 Daily Usage Scenarios

We also study the accounting discrepancy for seven users (university students), who have data plans with two US carriers. We record the data usage observed by the mobile phone and the one charged by the operator for two weeks (June 10 - 23, 2012), except that User 7 had only one-day record on June 22. Note that, this small-sample

user study may not well represent the common cases of daily usage for average users in our society. Instead, we intend to demonstrate how much the charging gap could be observed in reality, which depends on executed applications, usage patterns and locations. Among these users, the most popular applications are Web browsing and Gmail. Table 4.5 also shows other popular applications for each user, such as Gmap, Skype, YouTube, PPS, FaceBook, ebook reading, and games. Though data usage varies with users (from 47.1 MB to 900.2 MB) due to user behavior diversity, the volume gap is indeed small (< 1 MB usually) for most users in reality. Big volume gap is not commonly observed in practice due to the built-in control mechanisms in many applications and infrequent encounters of NS-zones. However, we still observe that Users 4 and 7 have experienced volume gaps as large as 5.3% and 7.2%, respectively. User 4 once watched VLC streaming three times during a day while staying and roaming around his office area with several NS-zones; User 7 watched video using YouTube or PPS on the train to/from New York City, where there is a long tunnel without signals. During the round trip, User 7 transmitted and received 72.4 MB, but was charged by the operator for 77.6 MB, with the gap being 7.2%.

User	Operator-I				Operator-II		
	1	2	3	4	5	6	7
Apps (excl. Web, Gmail)	Gmap	Stock Games	Skype, PPS FaceBook	YouTube PPS	ebook	-	YouTube PPS
V_{UE} (MB)	194.2	270.3	124.6	900.2	121.7	47.1	72.4
V_{OP} (MB)	192.6	270.0	129.4	948.4	120.9	47.3	77.6
Gap (MB)	-1.8	-0.3	4.8	48.2	-0.8	0.2	5.2
	-0.9%	-0.1%	3.9%	5.3%	-0.6%	0.4%	7.2%

Table 4.5: Volume gap for user studies during June 10-23, 2012 (User 7 had only one-day usage record on June 22, 2012).

4.4.4 Recommended Quick Fix

We now recommend quick fix to the overcharging issues. The fundamental problem is that, the 3G network takes an SGSN/GGSN-based charging approach, which only records the data volume traversing these intermediate steps on the end-to-end delivery path. They do not coordinate with end devices when making accounting decisions. Specifically, they never take explicit feedback from the end systems. Therefore, when failures occur over the downstream path after SGSN/GGSN, the charging system is not aware of the status of the mobile device, thus incurring overcharge. We now suggest three feedback mechanisms that help to remedy the problems. Note that our proposals are also applicable to packet drops due to weak signals, not only in NS-zones.

In the first proposal, the charging system takes explicit feedback regarding the status of the end device. For the downlink case, our solution can be implemented within the 3G infrastructure without interacting with the UE. We use the feedback from RNC to obtain more accurate data usage delivered to the UE device. We use the field “*RNC Unsent Data Volume*,” which records the data volume not delivered to UE, defined by the 3G standard [3GP08b]. RNC reports this record to SGSN, which computes the data volume successfully delivered to UE, i.e., $V_{succ} = V_{SGSN} - V_{RNC_unsent}$. It thus enables the operator to charge the user based on the data volume delivered to UE. This way, the huge gap (e.g., the 450 MB) can be eliminated.

We next fix problems with the session teardown in the absence of signals, where the charging can stop early to avoid overcharging. Our suggested solution is to deactivate the PDP context soon after the UE device cannot be reached. This can be implemented by the soft-state mechanism on the PDP context. We set a timer with the PDP context for UE, and the timer, as well as the PDP context, will be refreshed via the data delivery to/from UE. Note in three-hour DL-NS experiments, PDP context is not released in time because there is incoming traffic associated with it. This implies that the operator probably makes wrong decision that the PDP context should be kept alive. We thus

suggest refreshing the timer based on actual data delivery, or the paging of UE when the actual data usage is zero. This offers the 3G charging system an alternative feedback mechanism on the UE status, but may incur excessive control overhead.

The third feedback mechanism also helps to reduce overcharging. Whenever big data usage is generated, it should trigger an exception verification to check whether the charging makes sense or not. For example, when a data session lasts for an hour or produces about 100 MB data, the 3G core network should verify whether it is indeed normal charging practice. This can be done by sending a signaling message to RNC to query whether the UE status is normal. The RNC subsequently reports the UE status and facilitates SGSN/GGSN in its charging decision.

We note that Cisco has proposed overcharging protection for GGSN to be aware of lost radio coverage using the feedback of SGSN [Cis11]. It is to assess the device status due to lost coverage.

4.4.5 Progress Update and Carriers in Other Regions

We also run similar experiments with three other major carriers, one each in the US, China and Taiwan. All the observed results still hold in general. The minor difference is that, (1) three-hours charge for a 50 Kbps UDP flow in DL-NS experiments is observed for two major US operators, while at least 5.7 hour charge is observed for the third US operator, one hour charge is observed in China, and about 42 minutes occur in Taiwan; and (2) the maximal UDP source rate is smaller in China, e.g., the transition point in Figure 4.9(b) happens at 1 Mbps for the Chinese carrier. This is because the data rate supported over the wired Internet is smaller in China. We also conduct two-week usage studies for one user in China and two users in Taiwan; their observed gaps are negligible (<1 MB, within 0.5% error) because they mainly use them for Web, Gmail and SMS exchanges and the overall volume is small.

Our further work shows this accounting gap is not limited to wireless failures. It also

occurs when mobility (handoff) is involved [TPL13]. The blame still goes to the open-loop data accounting architecture. During a handoff, data suspension and discarded buffer in the original base station lead to the inconsistent view on the core network and the end device.

4.5 We Pay For What We Do Not Want

Following the second finding, we will show that users pay for what they do not want. Here, we describe the stealth spam attack, which is a new spam threat against mobile devices by exploiting the loopholes in current 3G/4G accounting system. It stealthily injects a large volume of spam data, which the mobile device may not be even aware of (e.g., after the mobile device already closes the data session on its side). This incurs extra payment on the mobile user. We explore its root cause and propose remedies.

4.5.1 Vulnerability Analysis

Stealth spam attack is different from conventional spam threats targeting mobile devices. Conventional spams include Email spam, SMS/MMS spam, junk image or video embedded in Web pages, etc. Users are typically aware of these annoying junk messages and may take actions to block them. In contrast, the stealth spam attack can be long lived, lasting several hours or more (observed in our experiments). The persistent spam session not only allows for the attacker to send a large volume of junk data, but also does it covertly. The users may be completely oblivious of such attacks.

Since cellular operators charge mobile users of incoming traffic, they are responsible to secure it. Ideally, they should only allow the incoming traffic with the consent from mobile users; arbitrary traffic from the external data networks should be blocked by the operators. In fact, the operators do provide security mechanisms to protect it. They deploy various middle-boxes such as network address translation (NAT) boxes

and firewalls [WQX11].

The deployment of NAT makes launching mobile spam attack a challenging task. Specifically, NAT offers two countermeasures against spam. First, it decouples network access from public reachability. The mobile UE is only allocated a private IP address (not reachable from the external network) when its bearer (i.e., PDP context) is activated. The UE is reachable from the public Internet only after NAT assigns it a translated IP address and a port number. This dynamic assignment only occurs when the UE initiates a data session (e.g., when starting a Google search or signing in Skype). Without the explicit activation from the UE side, data-charging operations never happen (as shown in the normal case of Figure 4.15). This tends to shield most conventional spam threats that send data to the UE via its IP address.

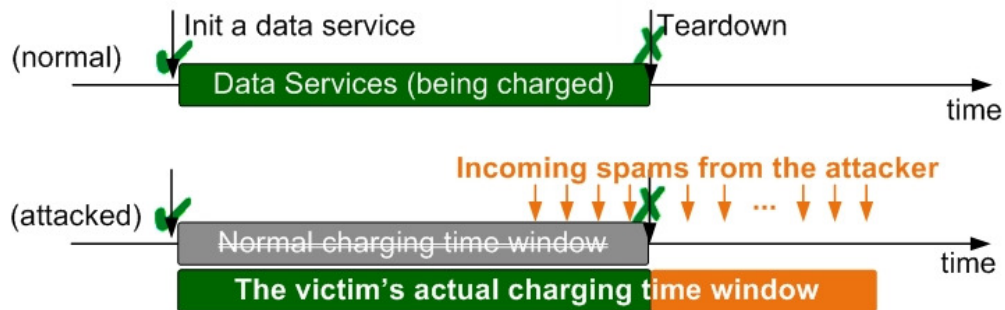


Figure 4.15: Illustration of the stealth spam attack.

As the second countermeasure, operator's NAT boxes only grant temporary permissions for the traffic traversing the cellular core network. They only allow for the traffic to pass through within a provisional time window when the data session is alive. In the normal scenario, the charging time window ends when the UE terminates this data service. For example, mobile Web browser may immediately send a TCP FIN message to close the TCP connection, once the Web page is downloaded. This way, only within the given time window, those hosts, which know the access information (i.e., the translated IP address and the port number), are able to inject traffic to the UE. This window-controlled access also helps to protect the UE from spam threat. In addition, firewalls deployed by operators can also filter out spam.

On the other hand, the loopholes in the current 3G/4G charging system, as well as in applications, also create opportunities for stealth spam attack. Our analysis and experiments show that, there exist two loopholes in the current charging system. The first loophole is that,

Data flow termination at the UE \neq charging termination at the operator.

The UE can terminate a data flow locally but *cannot* terminate the accounting of the flow. There exists inconsistency between the UE status and the operator's view on termination of a charging operation. When the user closes an application or an Internet service, (s)he thinks that the data flow is about to release and no more incoming traffic is allowed. However, the operator may view differently: This flow does not terminate as long as incoming packets belonging to this flow still arrive. The current 3G charging takes the operator's view. Therefore, charging can last much longer than expected. This occurs when the attacker starts this incoming spam before the normal teardown by the UE (shown in Figure 4.15)). We further find out that operators are unable to effectively stop data charging even when the UE explicitly sends teardown signals (e.g., in TCP). It is even worse for those UDP-based data service. The charging can last even longer once the spam starts; there is no sign for it to stop based on our experimental observation.

The second loophole is that,

Initial authentication \neq authentication during the whole data process.

All the authentication operations are performed at the start of the data flow (or when establishing the PDP context), but not when closing a flow. Therefore, the current charging procedure secures the initialization of the flow but not the whole process. Specifically, it cannot protect the data flow in the teardown process. The current design works for voice calls but not for data. Packet-switched IP data forwarding can push packets along different paths to reach the victim UE without prior consent, different from the circuit-switched fixed route for voice calls.




Figure 4.16: Two steps to launch the stealth spam attack.

With these loopholes, stealth spam attack can be launched. Figure 4.16 shows two typical steps to launch this attack: trap and spam. First, it traps the UE to obtain its confidential access information and flow permission to traverse the CN. The second step is to send junk packets. In the following, we describe how to implement them in several example scenarios and examine how badly it may hurt the victim.

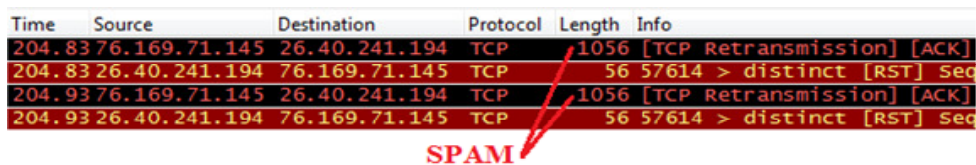
4.5.2 Spam Attack in TCP-based Services

We now describe how spam attack poses threats to those TCP-based services. Since TCP is a stateful protocol, we expect the spam to stop early once the UE application closes its TCP connection. Take Web browsing as an example. Once the Web page is fully retrieved, the Web browser may send a TCP FIN signal to the Web server and closes this TCP connection. Even though the Web server is malicious, the timeout mechanism also helps the UE to close this connection. The timer is typically set from tens of seconds to several minutes. However, our study has confirmed that the current charging practice contains loopholes. The operator may not stop charging, even when they can learn that the connection is closed by the UE.

In our experiments, we deploy a Web server as the attacker and modify its used TCP protocol. The spam attack starts when the UE clicks a malicious Web link and setups a TCP connection with the attacker. In the modified TCP, the normal TCP connection teardown procedure is disabled. This TCP will never send FIN or FIN-ACK signals like a normal TCP, upon receiving the teardown request from the UE. Once the UE

is connected, the attacker immediately sends junk packets at a fixed rate for a given duration. To enable fixed-rate testing, we also disable TCP congestion control.

We first run experiments using various source rates for five minutes. Figure 4.18 plots the data volume increase due to this attack in both networks. It is observed that, as the incoming source rate grows beyond one threshold (about 400Kbps for Operator-I, 200Kbps for Operator-II), the attack seems to be blocked by the operator. The higher the source rate, the earlier the attack is blocked. For example, the spam is blocked in 24.7 seconds when the incoming rate reaches 1 Mbps for Operator-I while it gets blocked in 2 minutes for those attack at the source rates from 300 Kbps to 1 Mbps for Operator-II. This result implies that, operators do offer certain protection mechanism (e.g., blocking the TCP connection if it is too fast). However, these protection policies are operator specific. We find that, Operator-I may block the access to any data service while Operator-II only blocks this specific data service. We also observe that, the charging time window is not determined by the TCP connection status. When the UE closes this TCP connection, it sends TCP-RESET signals upon receiving spam packets. Figure 4.17 shows the Wireshark trace at the victim; TCP-RESET signals indicate that the UE aborts the connection and the 1056byte-packets are spam data units. The trace shows that, operators still allow the delivery of those spam packets and charge mobile users even when the UE TCP connection is closed.




Time	Source	Destination	Protocol	Length	Info
204.83	76.169.71.145	26.40.241.194	TCP	1056	[TCP Retransmission] [ACK]
204.83	26.40.241.194	76.169.71.145	TCP	56	57614 > distinct [RST] Seq
204.93	76.169.71.145	26.40.241.194	TCP	1056	[TCP Retransmission] [ACK]
204.93	26.40.241.194	76.169.71.145	TCP	56	57614 > distinct [RST] Seq


SPAM

Figure 4.17: Wireshark traces at the victim after the UE tear downs the TCP connection.

We also test this attack at low source rate (150 Kbps) for various durations. Figure 4.19 plots the data volume increase in both operators. The low-rate attack can easily bypass the security check implemented by both operators. The attack can last for two hours; there is no sign to end during our experiments. The data volume incurred by this




(a) Operator-I




(b) Operator-II

Figure 4.18: Data volume caused by TCP-based stealth spam attacks at various source rates.



(a) Operator-I



(b) Operator-II

Figure 4.19: Data volume caused by TCP-based stealth spam attacks for various durations.

attack has exceeded 100 MB. There is no sign of limits.

4.5.3 Spam Attack in UDP-based Services


We now describe stealth spam in UDP-based services. Since UDP is connectionless, it is even harder to decide when the UDP-based service ends and when the charging operation ends accordingly. The bad news is that, there is no clear protection mechanism for UDP-based service, while the operators at least use sort of abnormality-check for TCP-based sessions. The malicious attacker can launch stealth spam in UDP-based services by trapping the victim to open a UDP connection with itself. It may not be

popular to use a malicious link to open UDP connection, we introduce to use two popular applications (e.g. VoIP and video streaming) to trap the victim and leak the access information.


Spam attack from your buddies We demonstrate that the attacker can use VoIP service, including Skype and Google Talk, to construct the stealth spam. We use Skype as the example application. Skype is a globally used VoIP service and allows users to communicate with peers via voice, video, and instant messaging over the Internet [Sky]. Skype allows the buddies to communicate directly. A buddy on Skype has the chance to directly connect to the victim device without extra authentication.

The first step to launch this attack is still to obtain the victim's confidential access information (i.e., translated IP address and port number) and its permission for this flow to traverse cellular networks (in the *trap* step). To this end, the attacker starts to make a call to this victim when he gets online using mobile phones. The attacker hangs up before the victim accepts the call, or even before the call rings at the victim side. This way, the victim may not be even unaware of this attempted call. During this process, the victim's Skype client performs two operations. First, it sends its access information to the attacker, which is proved in the attacker's Wireshark trace. In the meantime, it automatically notifies the operator that it accepts this flow, which subsequently grants the traffic flow from the spammer to traverse cellular networks. In the *spam* step, the attacker just sends junk UDP packets. The attacker can confirm that the victim is indeed a mobile user, based on the victim's translated IP address given by NAT. The operator-owned IP address block is readily known in advance. The spammer can also pick up the victim and launch operator-specific attacks.

We run experiments to validate this attack and verify whether extra checking mechanisms exist. We also vary the attack durations and incoming source rates in the tests. Figure 4.20 plots the overcharged volume versus different source rates during the five-minute Skype spam attack. The charging volume increase is in proportion to source rates. It implies that, operators do not enforce any security mechanism for UDP-based




(a) Operator-I




(b) Operator-II

Figure 4.20: Data volume caused by UDP-based (Skype) stealth spam attacks at various source rates.



(a) Operator-I



(b) Operator-II

Figure 4.21: Data volume caused by UDP-based (Skype) stealth spam attacks for various durations.

services. The spam volume can consequently grow much larger. We also make an interesting observation. In Operator-I, even though these packets may not be actually delivered to the UE (e.g., when the weak radio link cannot afford high-rate source), they are still charged by the operator. It shows that the operator might charge the mobile users based on the volume that arrive at them, not the one that they successfully delivery to the UE. Figure 4.21 plots the data volume caused by Skype stealth spam for various durations, with the source rate being 50kbps. It shows that the overcharge volume grows in proportion to the spam duration. There is no sign to end even when the attack has already lasted 24 hours for Operator-I (the overcharge volume reaches 500+ MB) during our experiments.

We also note that, the attack is still ongoing even after the victim signs out from Skype. The Wireshark trace at the victim side (see Figure 4.22) indicates that, spam packets still arrive at the UE and are charged by the operator after the UE logs out Skype. In the trace, the message of ICMP Port Unreachable shows that the UE has closed this application port after Skype logout.

In addition to Skype, this spam can be launched via Google Talk. The attacker also makes a call before the victim accepts it to trap the mobile user. The performance is similar; we omit it due to lack of space. Note that, the Skype/GTalk-based attack is a result of both 3G charging system vulnerability and Skype/GTalk implementation. The operator exposes the vulnerabilities at the first place, which still charges incoming spam packets that mobile application do not accept. The root cause is still that there is no feedback mechanism in the 3G charging system to tear down suspicious or malicious flows for mobile users. The Skype or other VoIP implementation (to release access information without explicit user confirmation) is exploited to mount this attack. Once you accept invitations from strangers or your buddies are compromised, you are vulnerable to this overcharging attack whenever you go online.

Time	Source	Destination	Protocol	Length	Info
173.98	131.179.210.21	26.41.147.68	UDP	1250	Sour: 27493 Dest: 44447
173.98	26.41.147.68	131.179.210.21	UDP	55	Sour: 44447 Dest: 27493
173.98	26.41.147.68	131.179.210.21	UDP	48	Sour: 44447 Dest: 27493
174.78	131.179.210.21	26.41.147.68	UDP	1250	Sour: 27493 Dest: 44447
174.78	26.41.147.68	131.179.210.21	ICMP	592	Destination unreachable
174.78	131.179.210.21	26.41.147.68	UDP	1250	Sour: 27493 Dest: 44447
174.99	131.179.210.21	26.41.147.68	UDP	1250	Sour: 27493 Dest: 44447
175.08	131.179.210.21	26.41.147.68	UDP	1250	Sour: 27493 Dest: 44447

SPAM
User has signed out

Figure 4.22: Wireshark traces at the victim after log out from Skype.

Spam attack in video streaming Other channels exist to launch stealth spam attack in UDP-based services. Video streaming is another example. To trap the victim, the attacker can create a malicious link to redirect Web-browsing operations to start a realtime video streaming. For example, the victim may click one phishing link which redirects the victim's browser to:

`rtsp://*.*.1.204:554/trackID=5,`

where RTSP is a network protocol to support video streaming [RFC98]. Once the link is clicked, the victim automatically starts a new RTSP (over UDP) session running on port 554 and releases its confidential access information to the attacker. Once completed, the attacker blasts spam packets. We implement this attack and test it. We find that it performs similarly to Skype spam attack since both run on the top of UDP. We omit it due to lack of space. In both cases, UDP-based spam can inject an arbitrarily large volume of traffic and force the UE to pay more.

In summary, we have demonstrated that the stealthy spam attack is a real threat to mobile users. The attack is rooted in the inherent loopholes in the current charging architecture. Unless these loopholes are fixed, mobile users may always be victims when the stealth attack or more sophisticated attacks built on it are launched. On the other hand, the good news for mobile users is that, there is no obvious and strong incentive for attackers to launch such attacks now. Attackers cannot have immediate gains for themselves, unless an ill-intentioned operator contracts hackers to attack its own users for larger revenue gain or attack users in its competitor's network for unexpected user complaints, or a disgruntled attacker uses it to incur large monetary loss against his adversary. However, we quickly admit that incentive is an independent and interesting topic to study. Attackers may come up with unexpected incentives to launch more sophisticated attacks in this category in the future.

4.5.4 Root Cause

The fundamental problem underlying the stealth spam attack is that, there is no feedback mechanism from the UE to the carrier's charging system. So the operator cannot block unwelcome traffic based on the UE's feedback. This is an inherent design limitation in the current 3G/4G accounting system.

There is a second blame. The IP-based push model makes spam attack easy. IP for-

warding can send data packets to anyone without prior consent. Security mechanisms offered via NAT and firewall fail to work during the data delivery process once the service flow starts. Consequently, it cannot shield incoming spam data when malicious hackers hijack the flow or when the victim later finds that (s)he is trapped. As aforementioned, packet-switched networks are unable to offer synchronized and consistent views among all the involved equipments. Even though the mobile user is able to terminate the malicious (or suspicious) service on its application layer locally, it cannot terminate the accounting operations done at the carrier side. Operators decide on what packets are accounted using their own rules, lacking a feedback mechanism to allow the mobile user to explicitly express what packets are wanted or unwanted. Therefore, malicious attackers can inject spam packets and deceive the carriers to charge the mobile user for data volume larger than requested.

This attack again demonstrate the design flaw rooted in the 3G/4G accounting architecture. It defines how the accounting is performed and makes mobile users vulnerable to over accounting risks.

4.5.5 Recommended Quick Fix

The open-loop accounting is an inherent design limitation in the current 3G/4G charging system. Given IP-pushing model, spam attacks can send arbitrary packets to the mobile victims without their consent. Given the current architecture weakness, a viable accounting system must have the following three components: (1) The mobile user himself must be aware of such potential attacks and apply precaution measures. He can simply limit the size of any automatic downloaded data (such as email fetching); (2) The UE must be able to detect unwanted traffic and send feedback. The current protocols at the network layer and the transport layer are designed with such feedback. However, many applications ignore unwanted data (e.g., Skype does so) in general. This has to be fixed to make them suitable to run over a metered charging service; (3)

The carriers must take feedback from the UE to stop unwanted traffic.

Specifically, regarding the feedback mechanisms from the UE, we propose three solution options: *implicit-block*, *explicit-allow* and *explicit-stop*. The *implicit-block* solution is to enforce the CN components such as GGSN and NAT boxes. It uses implicit hints from the UE to justify whether the ongoing traffic is welcome or not. Once the traffic is unwanted, the CN blocks this flow and stops charging. The key issue is what messages can serve as hints on whether the UE's data packets are still wanted or not. For TCP-based service, TCP-RESET messages are sent from the UE if the the corresponding TCP connection is torn down earlier by the UE. Our study shows that, mobile Web browsers start to send TCP RESET messages upon receiving unintended TCP packets one-minute after they send the FIN signals. In case of UDP-based service, the UE responds a *ICMP Port Unreachable* message to the external sender upon receiving UDP packets on those closed ports. Therefore, messages of TCP-RESET and ICMP-Port-Unreachable can serve as the hints for the CN. To make correct decision, the CN can further exchange this information with the UE and seek confirmation from the mobile user. Using these implicit feedbacks, the CN should effectively disable the suspicious flows delivered to the UE. A downside of this solution is that, it takes effects only if the UE explicitly tear downs the service (e.g., quitting an application, terminating a TCP connection).

In the *explicit-allow* remedy, the UE explicitly specifies which packets are anticipated by adding/modifying the Packet Filters of Traffic Flow Template (TFT) associated with its PDP context. It can be done using *MS-Initiated PDP Context Modification Procedure* defined by 3GPP [3GP06a]. The attributes of the packet filter include [3GP06a]: (1) remote address; (2) local address; (3) protocol number, i.e., IPv4; (4) local port range; and (5) remote port range, etc. By adding packet filters, the UEs may not suffer from large spam attack when they are trapped or cheated to receive unexpected packets. One possible downside is that, it requires the UE to be fully aware of what it intends to send/receive. It requires detailed domain knowledge on various applications

and services.

The *explicit-stop* solution is to provide explicit feedback from the UE to the carrier when closing some data services. Once the phone detects that there exists any malicious or suspicious flow, it immediately reports to the core network and asks to block such a flow. The spam flow can be detected by mobile anti-malware software, or identified by mobile applications or systems software (e.g., an exception is issued when the application layer or a lower layer in the protocol stack discards a large number of packets). Malicious attackers can also be detected through the collaboration of many phones [CWY07]. This solution framework is flexible enough to integrate with different detection options. It also allows for the UE to stop data charging at any time, even when the UE was cheated or unaware of the attack at the start of the service. Its downside is that, current 3G/4G standards do not offer such mechanisms.

4.6 Gray Areas in Data Accounting

We now describe accounting cases in gray areas, where the users may be charged differently by the operators, compared with the actual data usage perceived at end hosts. However, there is no simple, accepted rule in these cases. In addition to the above three findings, we will show that the users may be charged for wrong or careless uplink operations; the users may be charged due to the middlebox deployment and Internet traffic congestion on mobile data charging. We finally assess accounting of application overhead.

4.6.1 UDP Uplink to a Nonexistent Host

The worst uplink case is to use UDP packets to a nonexistent host (i.e., no packets can be successfully delivered). Our tests show that both operators still charge every bit sent by UE. The root cause still lies in the SGSN/GGSN based charging architecture,

similar to the downlink case of Chapter 4.4. The good news is that this scenario is not very common unless the device is hijacked.

We also test UDP uplink traffic to our server under various wireless environments. It turns out that, there is no (obvious) gap between the data volume arriving at the receiver (i.e., our server) and the volume charged by the operator. However, our UE traces show that the UE does retransmit data over the wireless link (particularly in the WR-zone/WS-zone), i.e., $V_{UE} > V_{OP}$. Note that, these retransmitted data over the wireless link will not be observed by SGSN/GGSN, thus incurring no extra charge beyond those volume perceived by the receiver.

We now consider TCP uplink case. Interestingly, we find out that, TCP acts the opposite role to the one in the downlink case. We discover that TCP increases the charged volume in some cases. Specifically, we discuss two cases with the middlebox deployment and Internet traffic congestion.

4.6.2 Effect of Middle-boxes

Middleboxes (e.g., proxy servers, CDN servers, NAT boxes, and firewalls) can be deployed inside 3G/4G networks for performance enhancement or extra service [WQX11]. Indeed, our study confirms that proxy servers are placed in 3G networks. We find out that, Operator-I deploys proxy servers to handle popular applications, including HTTP and FTP. Consequently, the end-to-end data session between the UE and the server is split into two segments, one between the UE and the proxy, the other between the proxy and the HTTP/FTP server. We now assess the impact of such session splits due to middleboxes on charging. In the worst case, the UE interacts with the proxy rather than its intended server, and incurs overcharging. The user is charged though (s)he never receives any service! This can be illustrated by the following experiment.

We let the mobile phone connect to a non-existent host (i.e., an unallocated IP address) and upload a 1 MB file using TCP. Since the host does not exist, it should

stop early (e.g., after several TCP SYN requests). To our surprise, we discover that, the data sessions for HTTP (80, 8080) and FTP (21) last much longer than expected. The delivered data volume for HTTP and FTP reaches 300 KB and 130 KB, respectively, in Operator-I. We also test with other popular applications, e.g., HTTPS(443) and SMTP(25), and unknown ports. Figure 4.6.2 plots V_{OP} (the same as V_{UE}) using different port numbers for both carriers. In Operator-I, we examined the TCP traces collected at the UE, and observed that those sent packets be acknowledged by TCP, though the IP address on the server side is nonexistent. This indicates that at least a middlebox has been deployed on the delivery path, which responds to UE requests on behalf of the server. The bad news is that, Operator-I also imposes charges though such HTTP/FTP data never reach the server side! We note that, the proxy is operator dependent and application specific. Other popular applications do not observe such proxy servers. Operator-II does not seem to have deployed such middleboxes in its core network even for HTTP/FTP.




Figure 4.23: Results when connecting to nonexistent hosts using different TCP ports

4.6.3 Packet Drop over the Internet

Charging discrepancy exists when packets are dropped over the wired Internet, which is outside the cellular core network. We illustrate the scenario in Figure 4.6.3. The UE sends 200 KB data to the Internet server via the cellular network, but packets are dropped right before reaching the server. Note that the Internet only offers best-effort service, so IP packet drops can be common, e.g., upon network congestion or malfunctioning routers. In this case, packets have been through the cellular network and recorded by SGSN/GGSN for charging. In our experiment, we upload an image file to our server from the mobile phone. We vary the packet-drop rate from 0% to 40% before reaching the TCP receiving end; we use it to emulate packet loss on the Internet. Since TCP will retransmit these lost packets from the UE, they incur additional charging volume at the carrier. We plot the charging gap (between the server and the operator) versus the drop percentage in Figure 4.6.3. The figure shows that, the charging volume increases almost in proportion to the drop rate. For operators, this makes perfect sense since the 3G network does deliver those packets. However, end users never receive those dropped data. The same phenomenon occurs for uplink UDP transmissions.

The question we raise is if the users should pay for extra bits? Here, we notice that Internet is beyond the control of cellular networks and there is no way to blame cellular networks for extra costs due to unreliable transmission that happens in the Internet. However, the UE does spend more to obtain the same kind of service. It is rooted in the difference between cellular-network charging system and the one for the Internet. The current Internet service providers choose unlimited data plans and we are not charged based on the traffic volume. Therefore, it is not a big deal if unreliable transmission happens and traffic volume increases. However, it does matter in cellular network scenarios because the mobile bills are based on the traffic volume. Their charging is isolated while the data delivery is not separated. Since reliable delivery is not guaranteed, volume-based charging is reasonable? Since the Internet and cellu-

lar networks affect each other, should their charging solutions be coordinated? These questions should be carefully addressed to design mobile data accounting system.




Figure 4.24: Results when experiencing unreliable Internet packet delivery.

4.6.4 Charging for Application Signaling

Motivated by the previous DNS study, we also examine whether application signaling messages are also charged by 3G accounting and how much percentage they contribute when charged. Our findings show that, both 3G operators charge the signaling data (e.g., ICMP, SIP and RTSP) and protocol overhead. However, the actual signaling cost varies a lot across different applications. We consider three interesting cases.

FTP Signaling Channel FTP uses two separate TCP sessions, with port 21 for command signaling and port 20 for data transfer. We conduct experiments to send messages mainly over the signaling channel (e.g., list a remote folder with 1 or 50 files), as well as data transfer (downloading one or ten 1MB files).

HTTP Redirect and Invalid Links We run HTTP redirect cases, where the web page is redirected once or 15 times to reach the final content. We also access a web page with one invalid HTTP link. In the invalid link case, we access a web page that has one or 50 invalid image links.

Email and IM We tested Yahoo Mail and Skype for Email and IM applications, re-

spectively. We send a small or large email, login/logout skype, or remain idle for 10 minutes in Skype.

The results are shown in Table 4.6. We make three observations. First, the application signaling messages, including FTP control commands, are indeed charged. Second, the signaling overhead percentage is particularly large when the content size is small, e.g., FTP listing, HTTP redirection. Command messages may only be a small percentage in the operator’s charging volume, compared with protocol overheads (see FTP signaling); Third, signaling and protocol overhead do incur hidden costs (not perceived by average users). Note that, in the HTTP case, those invalid links or those redirects are never the content requested by the user. It explains why Alice is charged by an invalid click without accessing real content in the motivative example.

Application	Test	Content (KB)	OP (KB)	Gap	Gap/OP (%)
FTP	Listing (1)	0.06	2.97	2.91	97.9
	Listing (50)	3.32	7.28	3.96	54.4
	Downloading (S)	1024	1190.5	166.5	14
	Downloading (L)	10240	10858.0	618.0	5.7
HTTP	Redirect (1)	0.05	1.9	1.85	97.3
	Redirect (15)	0.05	15.1	15.05	99.7
	Invalid (S)	0.13	2.05	1.92	93.6
	Invalid (L)	2.56	12	9.44	78.7
Email	Send (S)	0.02	13.0	12.98	99.8
	Send (L)	223.6	250.98	27.38	10.9
Skype	login/out	0	50.07	50.07	100
	idle (10mins)	0	5.05	5.05	100

Table 4.6: Signaling overhead of popular applications.

In a broader view, the above study makes us contemplate on who should pay for what. We have seen that free applications may raise more overhead for advertisements, while the paid ones may not. In recent years, 3G/4G operators have been making effort to evolve from “dumb-bit-pipe owners” to “content/service providers”. In the role-switching process, the charging policy may also evolve towards more content based;

lots of interesting issues and debates may arise.

4.7 Discussion and Conclusion

We conduct experiments on operational 3G networks to study mobile data accounting architecture and practice. We have discovered loopholes and showcased simple attacks, which are validated by experiments over two operational 3G networks. We notice that, the most optimistic view will claim that, the problem is not too bad, so we do not need to fix it; the built-in control mechanisms in TCP and applications at the end devices help to mitigate the damage. However, we believe that there are fundamental technical problems beneath these engineering missteps.

We believe that the open-loop accounting is an architecture flaw, since the 3G/4G core network has no means to learn what data is actually needed/perceived by the user. As a result, the accounting system still records data volume that are not delivered to mobile users, or spam data volume for users even after the user has terminated its data service. This calls for new feedback mechanisms between the end device and the core network for collaborative and consistent accounting actions. It inspires us to rethink accounting architecture. In general, there are three classes of accounting architecture: the network-based one such as the current 3G system, the end-system-based approach, and the collaborative one between the network and end devices. Both the network-based and the end-system-based approaches have severe limitations. For the network-based, 3G charging system, we already observed that it should result in large accounting gap in the extreme cases. The fundamental problem is that, the network lacks coordination with end systems and makes the charging decision alone. This functions fine when everything goes well, but suffers when things go wrong. The built-in feedback loop in TCP and applications may help to certain extent. However, the concrete feedback mechanism and its operation accuracy are largely out of control to the accounting system. The charging system is not self-healing under failures and extreme conditions. On

the other hand, the other extreme of end-system-based accounting will not work either. There is no easy mechanism to regulate users so that they will not cheat. The verification process of user-reported results is also challenging. Therefore, this approach is unrealistic in practice.

Our study yields some insights. On the policy side, differential charging seems to be a popular practice for mobile data services. Given a metered charging system, people necessarily have incentives to exploit and abuse any transfer that is free. There is no simple, bullet-proof solution except eliminating the free service. In the more general problem setting, as long as differential charging exists among applications and services, attackers have incentives to abuse transfers that charge less. The free service simply exemplifies an extreme case. While the toll-free-data attack seems to be readily fixed, we believe that more fundamental issues need to be addressed in the long run. The current 3G/4G accounting architecture lacks proper validation and verification on the traversing traffic types and content, when offering differential charging for applications. The scalability of the associated security design also needs to be considered because of the increasing traffic diversity and volume, as well as the large user population. On the architecture side, the charging system records the data volume on behalf of users, but does not take any user feedback when making charging decisions. So the carrier cannot block unwelcome traffic by using feedback from users. The IP-based push model makes spam attack easier. Anyone can send to the UE without prior consent. On both fronts, really bad things can happen, somewhat unexpectedly.

Given the current architecture weakness, a dependable, usage-based charging system calls for concerted renovations among the network, the mobile device, and applications. The mobile user himself must be aware of such threats and apply precaution measures. The UE must be able to detect unwanted traffic and send feedback. Many applications lack such feedback mechanisms and simply ignore unwanted data, e.g., in the case of Skype. This must be fixed to make them suitable to run over a metered charging service. The operators must take feedback from the UE to stop unwanted traffic, and

such feedback has to be carefully validated. The network also needs appropriate traffic validation and verification when making differential charging decisions for different applications and services. This work describes our current effort along this direction. We hope our preliminary study will stimulate further research on this important topic from both academia and industry.

CHAPTER 5

Conclusion and Future Work

We first summarize our work and then share the lessons and insights we have learnt. Finally, we present the immediate future work and the outlook of network research for mobile.

5.1 Summary of Results

The Internet is going wireless and mobile. Two underlying driving forces have been the explosive growth of smartphones/tablets and the rapid deployment of 3G/4G infrastructure. In this thesis, we have obtained two main results on cellular networks.

Traffic-driven energy saving in 3G infrastructure In this work, we have explored the feasibility of traffic-driven energy savings for operational cellular networks. Our study is motivated by two factors. First, deployment of base stations (BSes) in 3G networks is dense and redundant to ensure universal coverage; consequently, the operational cost is also high and this cost will be eventually transferred to mobile users who have to pay more for data access. Therefore, reducing the energy cost of the 3G/4G infrastructure can benefit users. Second, energy-efficient design has long been an active

research area in mobile networks. However, current study on energy saving is mainly on the mobile device side, while the infrastructure side story is largely overlooked. The 3G infrastructure consumes 90% to 99% of overall network energy in current operations. A practical and effective solution is definitely needed to build a *green* 3G/4G network. Driven by both demands, we explore a new approach to green cellular infrastructure. The key idea is to leverage the traffic dynamics and rich redundancy in BS deployment. We find out that BSes are not energy proportional with respect to their carried traffic loads. Our measurement further shows that 3G traffic exhibits high fluctuation both in time and over space, thus incurring energy waste. We then profile BS traffic and approximate network-wide energy proportionality using non-load-adaptive BSes. The instrument is to leverage temporal-spatial traffic diversity and node deployment heterogeneity, and power off under-utilized BSes under light traffic. The resulting solution is a location-dependent, profile-based scheme, which yields up to 52.7% savings in dense city network, and 23.4% in a mid-sized city with sparse deployment. Our proposal is also not cost free. It attains significant savings on the infrastructure side at the cost of mild increase on the clients during off-peak hours.

In the design process, we trade performance increments for simplicity, in that we always retain simple operations rather than squeeze every bit of possible gains. Instead of taking the popular optimization-based approach, we seek to design practical schemes that will work in reality. In a broader scope, our solution explores to build energy-proportional 3G networks using legacy non-energy-proportional base stations.

Fidelity and insecurity of 3G/4G data accounting system Unlike the wired Internet, cellular networks have implemented usage-based charging, rather than the simpler flat-rate charging. Going down this path, the 3G/4G standards finalize the accounting architecture, yet leave enough freedom for operators to define their own charging policy. In this work, we examine the accounting architecture and practice in 3G/4G cellular networks from both robustness and security perspectives. To this end, we first conduct the fidelity study through measurements and analysis. We have identified the

architectural limitation and loopholes in industry practice on data charging over operational 3G/4G cellular networks. In 3G/4G networks, mobile applications are charged based on the usage traffic volume.

We conduct experiments to critically assess both this usage-based accounting architecture and application-specific charging policies by operators. We have found that, both generally work in common scenarios but may *go wrong* in the extreme cases: We are charged for what we never get, and we can get what we want for free. In one extreme case, we are charged for at least three hours and 450 MB or more data despite receiving no single bit. In another extreme case, we are able to transfer 200 MB or any amount we specify for free. The root causes lie in lack of both coordination between the charging system and the end device, and prudent policy enforcement by certain operators. We further study the problem from the security perspective, and have discovered two effective attacks. The “toll-free-data-access-attack” enables the attacker to access any data service for free. The “stealth-spam-attack” incurs any large traffic volume to the victim, while the victim may not be even aware of such spam traffic. Our experiments on three US carriers have confirmed the feasibility and simplicity of such attacks.

We also propose defense remedies. Given the current architecture weakness, a dependable, usage-based charging system calls for concerted renovations among the network, the mobile device, and applications. The mobile user himself must be aware of such threats and apply precaution measures. The UE must be able to detect unwanted traffic and send feedback. Many applications lack such feedback mechanisms and simply ignore unwanted data, e.g., in the case of Skype. This must be fixed to make them suitable to run over a metered charging service. The operators must take feedback from the UE to stop unwanted traffic, and such feedback has to be carefully validated. The network also needs appropriate traffic validation and verification when making differential charging decisions for different applications and services.

5.2 *Insights and Lessons*

Our thesis study offers several insights.

1. Migrating circuit-switched design to packet-switched networks At first glance, the data accounting issues seem to be engineering glitches, which are inevitable in large-scale systems. However, we believe that the root causes lie in the underlying technology fundamentals: conflicts in design for circuit-switched systems and solutions to packet-switched networks.

In 3G/4G networks, the accounting design for packet data is inherited from the legacy accounting system for cellular voice. Since voice support has used the *circuit-switched* paradigm in telecom systems, the SGSN/GGSN element-based charging has been working well. Once the virtual circuit is established for voice, any element on the circuit observes the same, consistent view on the communication behavior, including volume accounting. When migrating this design from the circuit-switched systems to support data in packet-switched networks, problems arise. Packet delivery does not occur on a virtual-circuit path in a packet-switched 3G data network, and no service guarantee is ensured but only best-effort service is available. In a best-effort service model, various packet delivery misbehaviors may occur, e.g., packet loss, prolonged delay, out-of-order delivery, to name a few. Consequently, when things go wrong outside the charging elements, the resulting data volume deviates from what is observed at end devices. The fundamental problem is that, the charging system mainly uses open-loop, but not closed-loop operations. This open-loop design works well for circuit-switched systems, since the circuit uses closed-loop feedback in its circuit establishment, ongoing maintenance, and teardown. However, it does not hold for packet-switched networks, which do not have the end-to-end feedback at the network layer. As a result, in the worst-case scenario (e.g., a typical open-loop UDP session), it makes accounting decisions alone without taking feedback from end devices. A user may thus be charged for

what he never receives, as shown by our extensive experiments.

Moreover, in packet-switched systems, policy is a double-edged sword. On one hand, it offers flexibility and is a good mechanism for operators. On the other hand, policy practice is also mistake prone. Policy operators have to take extra care to make policy enforcement complete. Otherwise, really bad things can happen under extreme conditions, somewhat unexpectedly. Consequently, as shown by our experiments, we may get what we want for free in worst-case scenarios.

2. Diversity is pervasive in cellular networks Our study so far has revealed that, diversity is pervasive in current wireless networking systems, spanning both in time and over space along multiple dimensions. The observed diversity includes changing traffic dynamics over time, heterogeneous node deployment density, traffic dynamics in spatial neighborhood, diversified communication range and node capacity, etc. Such multi-dimensional diversity not only leads to energy inefficiency in wireless networks, but only offers new opportunity for green networking design when used properly. Our work in thesis has shown that, the traffic dynamics spanned both in time and over space can be leveraged for energy saving. Energy efficiency can be achieved via exploiting traffic diversity and near-term stability both in time and over space, thus ensuring temporal-spatial multiplexing to save more energy.

The inherent diversity also sheds lights on exploring a novel, diversity driven approach to green wireless networking. The key point is to not suppress diversity through *averaging* or *filtering*, but to leverage it to improve energy efficiency. The overall design philosophy is that, in a large-scale networking system, diversity is inherent in nature and goes hand in hand with scaling. In reality, it exhibits many forms of heterogeneity in both communication infrastructure and user clients. On the user demand side, we have heterogeneous applications and traffic, and customized operation modes based on user preferences. On the service side, we have diversified communication operations (in terms of offered capacity in a given area, coverage range of a given infrastructure component, and various operation modes to adapt to different operating

conditions) and non-uniform infrastructure deployment. While the research community has investigated how to use diversity to improve communication performance, no explicit effort has been made to achieve higher degree of energy efficiency.

3. Reexamining the current 3G/4G security Our study also shows that current security mechanisms and solutions in 3G/4G networks are insufficient for mobile data applications.

The security protection mechanisms designed for cellular networks are following the old practice and operations of telecom systems, i.e., the system is still largely closed and end devices can use but will not directly access the infrastructure. Therefore, two main security mechanisms are in place. One is the hardware-based authentication that prevents unauthorized users from accessing the network, i.e., the SIM-card based authentication and authorization. The other is the firewall and NAT based solution designed for IP traffic. However, as the cellular networks move toward IP-based, more open system for mobile applications, these mechanisms are not sufficient. Part of the problem is due to the insufficient security in IP design, part of the problem is from the old circuit-switched cellular systems inherited from the telecom community.

Specifically, from the data charging perspective, three problems will arise. First, the Internet largely uses the flat-rate based charging, therefore, its push model in IP forwarding does not encounter real problems. Users can push IP packets to the receiver without the receiver's consent. This model would not work well with the metered accounting scheme in 3G/4G networks. The IP-based push model makes spam attack easier. Anyone can send to the UE without prior consent. Second, policy practice also brings new loopholes for security. On the policy side, differential charging seems to be a popular practice for mobile data services. Given a metered charging system, people necessarily have incentives to exploit and abuse any transfer that is free. There is no simple, bullet-proof solution except eliminating the free service. In the more general problem setting, as long as differential charging exists among applications and services, attackers have incentives to abuse transfers that charge less. The free service simply

exemplifies an extreme case. While the toll-free-data attack seems to be readily fixed, we believe that more fundamental issues need to be addressed in the long run. Third, on the architecture side, the current 3G/4G accounting architecture lacks proper validation and verification on the traversing traffic types and content, when offering differential charging for applications. The scalability of the associated security design also needs to be considered because of the increasing traffic diversity and volume, as well as the large user population. On the architecture side, the charging system records the data volume on behalf of users, but does not take any user feedback when making charging decisions. So the carrier cannot block unwelcome traffic by using feedback from users. Consequently, as confirmed by our experiments, victims may be charged for what they never anticipate, and attackers get data services they never pay.

4. Rich signaling on control plane While cellular networks have migrated from the circuit-switched design to the IP-based packet-switched system, they have inherited the rich, somewhat complex signaling design on the control plane. On one hand, simplifying the control-plane operation poses a valuable long-term research task. On the other hand, we can leverage the available signaling messages to refine the protocol design and architecture option, particularly from the resiliency standpoint. For example, if we indeed need to take feedback from the end-user device, how to obtain such feedback information poses a design issue. If it comes directly from the end device, then we cannot avoid the challenging problem of who to trust and what to trust. This is because users may cheat, so their information cannot be fully trustworthy. To address this issue, we can leverage the rich information collected on the cellular infrastructure, which offers such information feedback on the status of the end-user device. We thus avoid the cheating problems by users. To certain extent, how to use the available signaling and control-plane information at the infrastructure raises an interesting research direction for wireless networking research in the age of mobile applications.

5.3 Immediate Future Work

Along the line of this thesis work, several pieces of immediate future work can be done.

Adding resiliency into mobile accounting While we propose immediate solution fixes to the identified problems in data accounting, we plan to explore a solution from the technology fundamentals. Our solution needs the following three basic building blocks: (a) Cellular accounting architecture needs new mechanisms to enhance its security against attacks and reliability against failures; (b) User offers feedback hints to the infrastructure; (c) Adaptive schemes to manage diversified root causes for overcharging. On each part, we plan to investigate what mechanisms are necessary and how to design such mechanisms in the overall architecture, and what changes are needed on the 3G/4G standard documents.

Software tool for user accounting Along the accounting work reported here, we plan to design and implement a software tool that works on the Android platform, which detects, alerts and reacts charging misbehavior. Given the current architectural weakness in 3G/4G networks that cannot be addressed overnight, we propose client-based solutions that can be implemented with the current technology. It records the traffic usage volume over the 3G/4G NIC on the mobile device. In the meantime, it checks the online usage record kept by carriers whenever suspicious behaviors are detected from the client device side, and compares the usage volume. When visible overcharging gap is observed, it initiates termination procedures to involuntarily tear down the suspicious data session. To reduce the potential communication signaling, it also uses the overcharging map, which is pre-fetched to the mobile device (e.g., when free WiFi access is available), to minimize runtime detection and probing frequency.

Furthermore, we design a simple online database that offers carrier-specific, overcharging map spanning geographic regions. The charging map records areas where suspicious billing is found, similar to the coverage map that logs accessibility and con-

nectivity for smartphone users. This online database archives the overcharging reports sent by users in certain regions for a given operator in the past. It exploits the user community effort to sketch out the overcharging map in given regions. It enables users to download and look up at any time, and to take precautionary measures before running mobile applications.

Diversity-driven energy saving We plan to further pursue the direction to green wireless by exploiting diversity inherent in the systems. We seek to approximate an EP System using Non-EP components and devise simple yet effective solutions within the standards framework. While our current work has focused on exploiting traffic dynamics, we plan to further explore other dimensions, particularly the different cell types and capacity heterogeneity for energy savings. The current cellular networks have deployed four types of cells over time, including macro-cell, micro-cell, pico-cell, and the emerging femtocell. These cells differ in their supported communication range, deployed density, and offered capacity. Such heterogeneous deployment offers another venue for ensuring higher degree of energy efficiency. We plan to examine how to dynamically power on/off a distinctive set of cell combinations to support different traffic demand by users. This further extends our current work that assumes homogeneous cell type.

5.4 Future Plan on Network Support for Mobile

The Internet is entering the era of “*mobile applications and services.*” In the meantime, mobile applications are transforming the way how people access and use network services. These facts pose both challenges and opportunities. As a result, both the US government (via NSF NeTS, SaTC, CSR, etc.) and industry have invested and continued to invest heavily in its R&D. Moreover, while big data becomes a top-priority research area, mobile big data holds great potential within the domain. To initiate mobile big data study, we need to build network and systems support, which is highly

dependable, adaptive, and extensible, for mobile applications.

Along this direction, I am interested in redesigning the cellular network infrastructure, to make it an open systems platform that is friendly to mobile devices and applications. The success story of the Internet offers a good example to start with. We can draw important lessons learned from the Internet design and valuable guidelines to operate the Internet systems. As a result, the cellular network will be simpler and easier to operate. For example, I believe that the complex inside-infrastructure operations (e.g., signaling and control messages and actions) can be greatly simplified. The data plane can also enhance its performance via cross-layer designs, e.g., leveraging the good design practice from the simpler WiFi systems. In the meantime, the success component in cellular networks, e.g., the wide-area mobility support can be retained. Finally, cellular networks also offer great opportunity to build security and resiliency into the system as a first principle. In essence, I believe that the current cellular infrastructure calls for an open architecture design that can support the mobile devices and applications as the first-class network citizen.

References

- [3GP06a] 3GPP. “TS23.060: GPRS; Service description; Stage 2.”, 2006.
- [3GP06b] 3GPP. “TS23.125: Overall High Level Functionality and Architecture Impacts of Flow Based Charging.”, Mar 2006.
- [3GP06c] 3GPP. “TS32.240:Telecommunication Management; Charging Management; Charging Architecture and Principles.”, Sep. 2006.
- [3GP07] 3GPP. “TS32.298:Telecommunication Management; Charging Management; Charging Data Record (CDR) Parameter Description.”, Sep. 2007.
- [3GP08a] 3GPP. “TS25.301: Radio Interface Protocol Architecture.”, Aug 2008.
- [3GP08b] 3GPP. “TS25.413: UTRAN Iu interface RANAP Signaling.”, Sep 2008.
- [3GP08c] 3GPP. “TS36.902: Self-configuring and Self-optimizing Network Use Cases and Solutions (V1.0.1).”, 2008.
- [3GP10] 3GPP. “TS 29.010: Base Station System and Base Station System Mobile-Services Switching Centre (BSS - MSC).”, 2010.
- [3GP11] 3GPP. “TS23.401: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access.”, Dec 2011.
- [3GP12a] 3GPP. “TS23.272: Circuit-Switched Fallback (CSFB) in EPS.”, 2012.
- [3GP12b] 3GPP. “TS45.001: Physical layer on the radio path; General description (v11.0.0).”, 2012.
- [4GA13] 4GAmerica. “Global 3G and 4G Deployment Status: HSPA, HSPA+ and LTE.”, 2013. <http://www.4gamericas.org/index.cfm?fuseaction=page&pageid=939>.
- [ABF08] David G. Andersen, Hari Balakrishnan, Nick Feamster, Teemu Koponen, Daekyeong Moon, and Scott Shenker. “Accountable internet protocol (aip).” In *SIGCOMM*, pp. 339–350, 2008.

- [ABG12] Pavan K. Athivarapu, Ranjita Bhagwan, Saikat Guha, Vishnu Navda, Ramachandran Ramjee, Dushyant Arora, Venkat N. Padmanabhan, and George Varghese. “RadioJockey: Mining Program Execution to Optimize Cellular Radio Usage.” In *ACM MOBICOM*, pp. 101–112, 2012.
- [ABI08] ABI Research. “Mobile networks go green—Minimizing power consumption and leveraging renewable energy.”, 2008.
- [ACR10] Hussam Abu-Libdeh, Paolo Costa, Antony Rowstron, Greg O’Shea, and Austin Donnelly. “Symbiotic Routing in Future Data Centers.” In *ACM SIGCOMM*, pp. 51–62, 2010.
- [AGM10] Mohammad Alizadeh, Albert Greenberg, David A. Maltz, Jitendra Padhye, Parveen Patel, Balaji Prabhakar, Sudipta Sengupta, and Murari Sridharan. “Data Center TCP (DCTCP).” In *ACM SIGCOMM*, pp. 63–74, 2010.
- [AK11] Siripuram Aditya and Sachin Katti. “FlexCast: Graceful Wireless Video Streaming.” In *ACM MOBICOM*, pp. 277–288, 2011.
- [ALV08] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. “A Scalable, Commodity Data Center Network Architecture.” In *ACM SIGCOMM*, 2008.
- [Ame10] 3G America. “Global 3G deployments UMTS HSPA HSPA+.”, Feb 2010.
- [AMR12] Myrto Arapinis, Loretta Mancini, Eike Ritter, Mark Ryan, Nico Golde, Kevin Redon, and Ravishankar Borgaonkar. “New Privacy Issues in Mobile Telephony: Fix and Verification.” In *ACM CCS*, pp. 205–216, 2012.
- [And] AndFTP. “AndFTP - LYSESOFT, v2.9.8.” <http://www.lysesoft.com>.
- [ARF10] Oliver Arnold, Fred Richter, Gerhard Fettweis, and Oliver Blume. “Power consumption modeling of different base station types in heterogeneous cellular networks.” In *Future Network and Mobile Summit’10*, 2010.
- [Aur91] Franz Aurenhammer. “Voronoi Diagrams: A Survey of a Fundamental Geometric Data Structure.” *ACM Comput. Surv.*, **23**(3):345–405, 1991.
- [BBV09] Niranjan Balasubramanian, Aruna Balasubramanian, and Arun Venkataramani. “Energy Consumption in Mobile Phones: a Measurement Study and Implications for Network Applications.” In *ACM IMC*, 2009.
- [BCJ12] Sourjya Bhaumik, Shoban Preeth Chandrabose, Manjunath Kashyap Jataprolu, Gautam Kumar, Anand Muralidhar, Paul Polakos, Vikram Srinivasan, and Thomas Woo. “CloudIQ: a Framework for Processing Base Stations in a Data Center.” In *ACM MOBICOM*, pp. 125–136, 2012.

- [BCM09] Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Rohan Murty, and Matt Welsh. “White Space Networking with Wi-Fi like Connectivity.” In *ACM SIGCOMM*, pp. 27–38, 2009.
- [BH07] Luiz André Barroso and Urs Hölzle. “The Case for Energy-Proportional Computing.” *Computer*, **40**(12):33–37, December 2007.
- [BMV10] Aruna Balasubramanian, Ratul Mahajan, and Arun Venkataramani. “Augmenting mobile 3G using WiFi.” In *ACM MobiSys*, pp. 209–222, 2010.
- [BNC10] Sourjya Bhaumik, Girija Narlikar, Subhendu Chattopadhyay, and Satish Kanugovi. “Breathe to stay cool: Adjusting cell sizes to reduce energy consumption.” In *Green Networking*, 2010.
- [BOL09] Biljana Badic, Timothy O’Farrell, Pavel Loskot, and Jianhua He. “Energy efficient radio access architectures for green radio: large versus small cell size deployment.” In *VTC’09*, 2009.
- [bs 07] “Worldwide Base Stations to Pass 3 Million by Year-End.”, 2007. <http://www.cellular-news.com/story/27002.php>.
- [bs 12] “Small Cells Outnumber Cellular Base Stations.”, Nov 2012. <http://www.telecoms.com/51947/small-cells-outnumber-cellular-base-stations/>.
- [Bus13] BusinessWire. “Apple’s App Store Marks Historic 50 Billionth Download.”, May 2013. <http://www.businesswire.com/news/home/20130516005403/en/Apple’s-App-Store-Marks-Historic-50-Billionth>.
- [Cai] Caida. <http://www.caida.org/research/traffic-analysis/tcpudpratio/>.
- [CBA10] Ionut Constandache, Xuan Bao, Martin Azizyan, and Romit Roy Choudhury. “Did You See Bob?: Human Localization using Mobile Phones.” In *ACM MOBICOM*, pp. 149–160, 2010.
- [CDE12] James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, J. J. Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaura, David Nagle, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Yasushi Saito, Michal Szymaniak, Christopher Taylor, Ruth Wang, and Dale Woodford. “Spanner: Google’s Globally-Distributed Database.” In *USENIX OSDI*, pp. 251–264, 2012.
- [CDG06] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E.

- Gruber. “Bigtable: A Distributed Storage System for Structured Data.” In *USENIX OSDI*, pp. 205–218, Berkeley, CA, USA, 2006.
- [CFG11] Erika Chin, Adrienne Porter Felt, Kate Greenwood, and David Wagner. “Analyzing Inter-Application Communication in Android.” In *ACM MOBISYS*, pp. 239–252, 2011.
- [cis] “Cisco ASR 5000 Series Serving GPRS Support Node Administration Guide.” [http://www.cisco.com/en/US/docs/wireless/asr_5000/12_0/OL-24828_SGSN_Admin.pdf](http://www.cisco.com/en/US/docs/wireless/asr_5000/12/_0/OL-24828_SGSN_Admin.pdf).
- [Cis11] Cisco. “Cisco visual networking index: Global mobile data traffic forecast update, 2010–2015.”, Feb 2011.
- [Cis13] Cisco. “Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012–2017.”, 2013. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html.
- [CK74] Vinton G. Cerf and Robert E. Khan. “A Protocol for Packet Network Intercommunication.” *IEEE Transactions on Communications*, **22**:637–648, 1974.
- [Cor13] IDC (International Data Corporation). “Worldwide Smart Connected Device Market Crossed 1 Billion Shipments in 2012.”, March 2013. <http://www.idc.com/getdoc.jsp?containerId=prUS24037713>.
- [CSB10] Jay Chen, Lakshmi Subramanian, and Eric Brewer. “SMS-based Web Search for Low-End Mobile Devices.” In *ACM MOBICOM*, pp. 125–136, 2010.
- [CSE93] Ron Cocchi, Scott Shenker, Deborah Estrin, and Lixia Zhang. “Pricing in Computer Networks: Motivation, Formulation, and Example.” *IEEE/ACM Transactions on Networking*, **1**:614–627, 1993.
- [CWY07] Jerry Cheng, Starsky H.Y. Wong, Hao Yang, and Songwu Lu. “Smart-Siren: virus detection and alert for smartphones.” In *ACM MobiSys*, 2007.
- [CZM11] Mosharaf Chowdhury, Matei Zaharia, Justin Ma, Michael I. Jordan, and Ion Stoica. “Managing Data Transfers in Computer Clusters with Orchestra.” In *SIGCOMM*, 2011.
- [DBM10] Katerina Dufkova, Milan Bjelica, Byongkwon Moon, Lukas Kencl, and Jean-Yves Le Boudec. “Energy Savings for Cellular Network with Evaluation of Impact on Data Traffic Performance.” In *European Wireless*, 2010.

- [DG04] Jeffrey Dean and Sanjay Ghemawat. “MapReduce: Simplified Data Processing on Large Clusters.” In *USENIX OSDI*, pp. 10–10, Berkeley, CA, USA, 2004.
- [dns] “DNSSEC.” <http://www.dnssec.net/>.
- [DNS11] Supratim Deb, Kanthi Nagaraj, and Vikram Srinivasan. “MOTA: Engineering an Operator Agnostic Mobile Service.” In *ACM MOBICOM*, pp. 133–144, 2011.
- [EL04] Tomas Edler and Susanne Lundber. “Energy Efficiency Enhancements in Radio Access Networks.” In *Ericsson Review*, 2004.
- [Eri] Ericsson. “RBS 3418 Product Description.”
- [Eri08] Ericsson. “Energy-saving solutions helping mobile operators meet commercial and sustainability goals worldwide.”, June 2008.
- [ETM05] William Enck, Patrick Traynor, Patrick McDaniel, and Thomas La Porta. “Exploiting Open Functionality in SMS-Capable Cellular Networks.” In *ACM CCS*, pp. 393–404, 2005.
- [Ezz05] Zoheir Ezziane. “Charging and Pricing Challenges for 3G Systems.” *IEEE Communications Surveys and Tutorials*, 7(1-4):58–68, 2005.
- [FBG11] Giordano Fusco, Milind Buddhikot, Himanshu Gupta, and Sivarama Venkatesan. “Finding green spots and turning the spectrum dial: Novel techniques for green mobile wireless networks.” In *DySPAN’11*, Aachen, Germany, 2011.
- [FFC11] Adrienne Porter Felt, Matthew Finifter, Erika Chin, Steve Hanna, and David Wagner. “A survey of mobile malware in the wild.” In *Proceedings of SPSM’11*, 2011.
- [FLM10] Hossein Falaki, Dimitrios Lymberopoulos, Ratul Mahajan, Srikanth Kandula, and Deborah Estrin. “A First Look at Traffic on Smartphones.” In *ACM IMC*, pp. 281–287, 2010.
- [frea] “Fast and Free Facebook Mobile Access with 0.facebook.com.” <https://www.facebook.com/blog/blog.php?post=391295167130>.
- [freb] “Free Fast Public DNS Servers List.” <http://theos.in/windows-xp/free-fast-public-dns-server-list/>.
- [frec] “Free GPRS Hack For Reliance Mobile.” http://www.megaleecher.net/Reliance_Internet_GPRS_Hack.

- [fred] “Free Gprs Mobile Tricks.” <http://darkwap.mobi/gprs-tricks/Free-Gprs-Mobile-Tricks>.
- [free] “Three PAYG Mobile Internet for FREE.” <http://www.digitalworldz.co.uk/226311-three-payg-mobile-internet.html>.
- [Fref] FreeProxy. <http://www.handcraftedsoftware.org/>.
- [FZ08] Gerhard P. Fettweis and E. Zimmermann. “ICT energy consumption—trends and challenges.” In *11th International Symposium on Wireless Personal Multimedia Communications*, Lapland, Finland, Sep 2008.
- [GGL03] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. “The Google File System.” In *ACM SOSP*, pp. 29–43, 2003.
- [GHJ09] Albert Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A. Maltz, Parveen Patel, and Sudipta Sengupta. “VL2: a Scalable and Flexible Data Center Network.” In *ACM SIGCOMM*, 2009.
- [GHM05] Albert Greenberg, Gisli Hjalmtysson, David A. Maltz, Andy Myers, Jennifer Rexford, Geoffrey Xie, Hong Yan, Jibin Zhan, and Hui Zhang. “A Clean Slate 4D Approach to Network Control and Management.” *SIGCOMM Comput. Commun. Rev.*, **35**(5):41–54, October 2005.
- [GK] Aditya Gudipati and Sachin Katti. “Strider: automatic rate adaptation and collision handling.” In *Proceedings of the SIGCOMM 2011 conference on SIGCOMM*, SIGCOMM ’11.
- [GKW13] Younghwan Go, Denis Foo Kune, Shinae Woo, Kyoungsoo Park, and Yongdae Kim. “Towards Accurate Accounting of Cellular Data for TCP Retransmission.” In *ACM HotMobile*, 2013.
- [GLL09] Chuanxiong Guo, Guohan Lu, Dan Li, Haitao Wu, Xuan Zhang, Yunfeng Shi, Chen Tian, Yongguang Zhang, and Songwu Lu. “BCube: a High Performance, Server-Centric Network Architecture for Modular Data Centers.” In *SIGCOMM*, 2009.
- [GLS12] Daniel B. Giffin, Amit Levy, Deian Stefan, David Terei, David Mazières, John C. Mitchell, and Alejandro Russo. “Hails: Protecting Data Privacy in Untrusted Web Applications.” In *USENIX OSDI*, pp. 47–60, 2012.
- [goo12] “Google Play Store hits 25 billion downloads.”, Sep 2012. http://news.cnet.com/8301-1023_3-57520324-93/google-play-store-hits-25-billion-downloads-launches-discount.

- [GRT10] Pradeep Kumar Gunda, Lenin Ravindranath, Chandramohan A. Thekkath, Yuan Yu, and Li Zhuang. “Nectar: Automatic Management of Data and Computation in Data Centers.” In *USENIX OSDI*, pp. 1–8. USENIX Association, 2010.
- [GS03] Maruti Gupta and Suresh Singh. “Greening of the Internet.” In *ACM SIGCOMM*, pp. 19–26, 2003.
- [HAB13] Kurtis Heimerl, Kashif Ali, Joshua Blumenstock, Brian Gawalt, and Eric Brewer. “Expanding Rural Cellular Networks with Virtual Coverage.” In *USENIX NSDI*, pp. 283–296, Berkeley, CA, USA, 2013.
- [HAD12] Dongsu Han, Ashok Anand, Fahad Dogar, Boyan Li, Hyeontaek Lim, Michel Machado, Arvind Mukundan, Wenfei Wu, Aditya Akella, David G. Andersen, John W. Byers, Srinivasan Seshan, and Peter Steenkiste. “XIA: Efficient Support for Evolvable Internetworking.” In *USENIX NSDI*, Berkeley, CA, USA, 2012.
- [HAR10] Andreas Haeberlen, Paarijaat Aditya, Rodrigo Rodrigues, and Peter Druschel. “Accountable Virtual Machines.” In *USENIX OSDI*, pp. 1–16, 2010.
- [HBB11] Ziaul Hasan, Hamidreza Boostanimehr, and Vijay K. Bhargava. “Green Cellular Networks: A Survey, Some Research Issues and Challenges.” *IEEE Communications Surveys and Tutorials*, **13**(4):524–540, 2011.
- [HGX11] Iztok Humar, Xiaohu Ge, Lin Xiang, Minh Jo, Min Chen, and Jing Zhang. “Rethinking energy efficiency models of cellular networks with embodied energy.” *IEEE Network*, March-April 2011.
- [HKP11] Daniel Halperin, Srikanth Kandula, Jitendra Padhye, Paramvir Bahl, and David Wetherall. “Augmenting Data Center Networks with Multi-Gigabit Wireless Links.” In *ACM SIGCOMM*, pp. 38–49, 2011.
- [HQG12] Junxian Huang, Feng Qian, Alexandre Gerber, Z. Morley Mao, Subhabrata Sen, and Oliver Spatscheck. “A Close Examination of Performance and Power Characteristics of 4G LTE Networks.” In *ACM MobiSys*, 2012.
- [HSJ12] Sangtae Ha, Soumya Sen, Carlee Joe-Wong, Youngbin Im, and Mung Chiang. “TUBE: Time-dependent Pricing for Mobile Data.” In *ACM SIGCOMM*, 2012.
- [HT07] Harri Holma and Antti Toskala. *WCDMA for UMTS: HSPA Evolution and LTE*. John Wiley & Sons, Inc., New York, NY, USA, 2007.
- [HT11] H. Holma and A. Toskala. *LTE for UMTS: Evolution to LTE-Advanced*. Wiley, 2011.

- [Hua] Huawei. “Efficient power amplifier: the trend for the development of Node B.”
- [Ins10] China Mobile Research Institute. “C-RAN: Road towards green radio access network.” In *White Paper, V1.0.0*, April 2010.
- [Iod] Iodine. <http://code.kryo.se/iodine/>.
- [ip] “IP Tunneling Through Nameservers.” <http://slashdot.org/story/00/09/10/2230242/ip-tunneling-through-nameservers>.
- [ITU13] ITU. “ICT Statistics.”, 2013. <http://www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx>.
- [JCK11] Mayank Jain, Jung Il Choi, Taemin Kim, Dinesh Bharadia, Siddharth Seth, Kannan Srinivasan, Philip Levis, Sachin Katti, and Prasun Sinha. “Practical, Real-Time, Full Duplex Wireless.” In *ACM MOBICOM*, pp. 301–312, 2011.
- [JST09] Van Jacobson, Diana K. Smetters, James D. Thornton, Michael F. Plass, Nicholas H. Briggs, and Rebecca L. Braynard. “Networking Named Content.” In *ACM CoNEXT*, pp. 1–12, 2009.
- [KCZ12] Taesoo Kim, Ramesh Chandra, and Nickolai Zeldovich. “Efficient Patch-based Auditing for Web Application Vulnerabilities.” In *USENIX OSDI*, pp. 193–206, 2012.
- [KKH12] Denis Foo Kune, John Koelndorfer, Nicholas Hopper, and Yongdae Kim. “Location Leaks on the GSM Air Interface.” In *NDSS*, 2012.
- [KLC12] Lorenzo Keller, Anh Le, Blerim Cici, Hulya Seferoglu, Christina Fragouli, and Athina Markopoulou. “MicroCast: Cooperative Video Streaming on Smartphones.” In *ACM MOBISYS*, pp. 57–70, 2012.
- [KP02] Michel Kouadio and Udo Pooch. “A Aaxonomy and Design Considerations for Internet Accounting.” *SIGCOMM Comput. Commun. Rev.*, **32**(5):39–48, November 2002.
- [LBW09] Patrick P. C. Lee, Tian Bu, and Thomas Woo. “On the Detection of Signaling DoS Attacks on 3G/WiMax Wireless Networks.” *Computer Networks*, **53**(15):2601–2616, October 2009.
- [Lea05] Neal Leavitt. “Mobile Phones: The Next Frontier for Hackers?” *IEEE Computer*, **38**(4):20–23, 2005.
- [LGY12] Hongbo Liu, Yu Gan, Jie Yang, Simon Sidhom, Yan Wang, Yingying Chen, and Fan Ye. “Push the Limit of WiFi based Localization for Smartphones.” In *ACM MOBICOM*, pp. 305–316, 2012.

- [LLY10] Kyunghan Lee, Joohyun Lee, Yung Yi, Injong Rhee, and Song Chong. “Mobile Data Offloading: How Much can WiFi deliver?” In *ACM CoNEXT*, pp. 26:1–26:12, 2010.
- [LSY11] Nikolaos Laoutaris, Michael Sirivianos, Xiaoyuan Yang, and Pablo Rodriguez. “Inter-Datacenter Bulk Transfers with Netstitcher.” In *ACM SIGCOMM*, pp. 74–85, 2011.
- [LWG13] Sangmin Lee, Edmund L. Wong, Deepak Goel, Mike Dahlin, and Vitaly Shmatikov. “Box: a Platform for Privacy-Preserving Apps.” In *USENIX NSDI*, pp. 501–514, 2013.
- [MAB08] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. “OpenFlow: Enabling Innovation in Campus Networks.” *SIGCOMM Comput. Commun. Rev.*, **38**(2):69–74, March 2008.
- [MCC09] Marco Ajmone Marsan, Luca Chiaraviglio, Delia Ciullo, and Michela Meo. “Optimal energy savings in cellular access networks.” In *GreenComm’09*, 2009.
- [Mis04] Ajay R. Mishra. *Fundamentals of cellular network planning and optimization: 2G/2.5G/3G...evolution to 4G*. Wiley, 2004.
- [MMT08] Xuehong Mao, Amine Maaref, and Koon Hoo Teo. “Adaptive Soft Frequency Reuse for Inter-Cell Interference Coordination in SC-FDMA Based 3GPP LTE Uplinks.” In *GLOBECOM*, pp. 4782–4787, 2008.
- [moba] “MobileMe.” <http://www.apple.com/mobileme/>.
- [Mobb] Opera Mobile. <http://www.opera.com/mobile/>.
- [Mot] Motorola. “Horizon 3G-n macro outdoor data sheet.”
- [MRS09] Min Mun, Sasank Reddy, Katie Shilton, Nathan Yau, Jeff Burke, Deborah Estrin, Mark Hansen, Eric Howard, Ruth West, and Péter Boda. “PEIR, the Personal Environmental Impact Report, as a Platform for Participatory Sensing Systems Research.” In *ACM MOBISYS*, pp. 55–68, 2009.
- [Nat12] Suman Nath. “ACE: Exploiting Correlation for Energy-Efficient and Continuous Context Sensing.” In *ACM MOBISYS*, pp. 29–42, 2012.
- [NEF12] Edmund B. Nightingale, Jeremy Elson, Jinliang Fan, Owen Hofmann, Jon Howell, and Yutaka Suzue. “Flat Datacenter Storage.” In *USENIX OSDI*, pp. 1–15, Berkeley, CA, USA, 2012.
- [NKN12] Karthik Nagaraj, Charles Killian, and Jennifer Neville. “Structured Comparative Analysis of Systems Logs to Diagnose Performance Problems.” In *USENIX NSDI*, pp. 26–26, 2012.

- [NST] NSTX. <http://thomer.com/howtos/nstx.html>.
- [NWG10] Zhisheng Niu, Yiqun Wu, Jie Gong, and Zexi Yang. “Cell zooming for cost-efficient green cellular networks.” *IEEE Commutation Magazine*, **48**:74–79, November 2010.
- [Odl01] Andrew Odlyzko. “Internet Pricing and the History of Communications.” *Computer Networks*, **36**:493–517, 2001.
- [OK10] Eunsung Oh and Bhaskar Krishnamachari. “Energy Savings through Dynamic Base Station Switching in Cellular Wireless Access Networks.” In *GLOBECOM*, pp. 1–5, 2010.
- [OKL11] Eunsung Oh, Bhaskar Krishnamachari, Xin Liu, and Zhisheng Niu. “Toward Dynamic Energy-Efficient Operation of Cellular Network Infrastructure.” *IEEE Communications Magazine*, **49**(6):56–61, 2011.
- [Pan12] Pratap Kumar Panigrahi. “Green Energy: A Perspective for Indian Rural Telecom.” *Journal of Green Engineering*, **2**, 2012.
- [PDH12] Pawan Prakash, Advait Dixit, Y. Charlie Hu, and Ramana Kompella. “The TCP Outcast Problem: Exposing Unfairness in Data Center Networks.” In *USENIX NSDI*, pp. 30–30, 2012.
- [PLT12] Chunyi Peng, Chi-Yu Li, Guan-Hua Tu, Songwu Lu, and Lixia Zhang. “Mobile Data Charging: New Attacks and Countermeasures.” In *ACM CCS*, 2012.
- [PPS] “PPS.” <http://www.pps.tv/>.
- [Pro] ProxyDroid. <https://play.google.com/store/apps/details?id=org.proxydroid>.
- [PTL12] Chunyi Peng, Guan hua Tu, Chi yu Li, and Songwu Lu. “Can We Pay for What We Get in 3G Data Access?” In *ACM MOBICOM*, 2012.
- [QBR11] Chuan Qin, Xuan Bao, Romit Roy Choudhury, and Srihari Nelakuditi. “TagSense: a Smartphone-based Approach to Automatic Image Tagging.” In *ACM MOBISYS*, pp. 1–14, 2011.
- [QM12] Zhiyun Qian and Z. Morley Mao. “Off-Path TCP Sequence Number Inference Attack-How Firewall Middleboxes Reduce Security.” In *IEEE S&P*, 2012.
- [QMX12] Zhiyun Qian, Z. Morley Mao, and Yinglian Xie. “Collaborative TCP Sequence Number Inference Attack: How to Crack Sequence Number under a Second.” In *ACM CCS*, pp. 593–604, 2012.

- [QQH12] Feng Qian, Kee Shen Quah, Junxian Huang, Jeffrey Erman, Alexandre Gerber, Zhuoqing Mao, Subhabrata Sen, and Oliver Spatscheck. “Web Caching on Smartphones: Ideal vs. Reality.” In *ACM MobiSys*, pp. 127–140, 2012.
- [QWG10a] Feng Qian, Zhaoguang Wang, Alexandre Gerber, Z. Morley Mao, Subhabrata Sen, and Oliver Spatscheck. “TOP: Tail Optimization Protocol For Cellular Radio Resource Allocation.” In *IEEE ICNP*, pp. 285–294, Washington, DC, USA, 2010.
- [QWG10b] Feng Qian, Zhaoguang Wang, Alexandre Gerber, Zhuoqing Morley Mao, Subhabrata Sen, and Oliver Spatscheck. “Characterizing Radio Resource Allocation for 3G Networks.” In *ACM IMC*, 2010.
- [QWG11] Feng Qian, Zhaoguang Wang, Alexandre Gerber, Zhuoqing Mao, Subhabrata Sen, and Oliver Spatscheck. “Profiling Resource Usage for Mobile Applications: a Cross-Layer Approach.” In *ACM MobiSys*, pp. 321–334, 2011.
- [QWX12] Zhiyun Qian, Zhaoguang Wang, Qiang Xu, Z. Morley Mao, Ming Zhang, and Yi-Min Wang. “You Can Run, but You Cannot Hide: Exposing Network Location for Targeted DoS Attacks in Cellular Networks.” In *NDSS*, 2012.
- [Rep11] Daily Mail Reporter. “Up to 20 million Americans ‘Overcharged’ by AT&T for Data Usage.”, 2011.
- [RFC98] “RFC2326: Real Time Streaming Protocol (RTSP).”, 1998.
- [RFC10] “RFC5966: DNS Transport over TCP - Implementation Requirements.”, 2010.
- [RFF09] Fred Richter, Albrecht J. Fehske, and Gerhard P. Fettweis. “Energy efficiency aspects of base station deployment strategies in cellular networks.” In *VTC’09 Fall*, 2009.
- [RFM10] Fred Richter, Albrecht J. Fehske, Patrick Marsch, and Gerhard P. Fettweis. “Traffic demand and energy efficiency in heterogeneous cellular mobile radio networks.” In *IEEE 71st VTC’10 Spring*, pp. 1–6, Taipei, Taiwan, May 2010.
- [RFR12] Mark Reitblatt, Nate Foster, Jennifer Rexford, Cole Schlesinger, and David Walker. “Abstractions for Network Update.” In *ACM SIGCOMM*, pp. 323–334, 2012.
- [RHS06] Fabio Ricciato, E Hasenleithner, Philipp Svoboda, and Wolfgang Fleischer. “On the Impact of Unwanted Traffic onto a 3G Network.” In *IEEE SECPeU*, 2006.

- [RLL12] Moo-Ryong Ra, Bin Liu, Tom F. La Porta, and Ramesh Govindan. “Medusa: a Programming Framework for Crowd-Sensing Applications.” In *ACM MobiSys*, pp. 337–350, 2012.
- [SHP12] Ankit Singla, Chi-Yao Hong, Lucian Popa, and P. Brighten Godfrey. “Jellyfish: Networking Data Centers Randomly.” In *USENIX NSDI*, pp. 17–17, 2012.
- [SJH12] Soumya Sen, Carlee Joe-Wong, Sangtae Ha, and Mung Chiang. “Pricing Data: A Look at Past Proposals, Current Plans, and Future Trends.” *CoRR*, **abs/1201.4197**, 2012.
- [Sky] Skype. <http://www.skype.com>.
- [SKY11] Kyuho Son, Hongseok Kim, Yung Yi, and Bhaskar Krishnamachari. “Base Station Operation and User Association Mechanisms for Energy-Delay Tradeoffs in Green Cellular Networks.” *IEEE Journal on Selected Areas in Communications (JSAC)*, **29**:1525–1536, 2011.
- [SNR09] Ashish Sharma, Vishnu Navda, Ramachandran Ramjee, Venkata N. Padmanabhan, and Elizabeth M. Belding. “Cool-Tether: Energy Efficient on-the-fly WiFi Hot-spots using Mobile Phones.” In *ACM CoNEXT*, pp. 109–120, 2009.
- [SNR10] Aaron Schulman, Vishnu Navda, Ramachandran Ramjee, Neil Spring, Pralhad Deshpande, and et.al. “Bartendr: A Practical Approach to Energy-Aware Cellular Data Scheduling.” In *ACM MOBICOM*, 2010.
- [SSL13] Hamed Soroush, Keen Sung, Erik Learned-Miller, Brian Neil Levine, and Marc Liberatore. “Disabling GPS is Not Enough: Cellular Location Leaks over the Internet.” In *PETS*, July 2013.
- [SSO08] Srinivas Shakkottai, R. Srikant, Asuman E. Ozdaglar, and Daron Acemoglu. “The Price of Simplicity.” *IEEE Journal on Selected Areas in Communications*, **26**(7):1269–1276, 2008.
- [Sta09] IEEE Standard. “IEEE 802.11n-2009 Amendment 5: Enhancements for Higher Throughput.”, 2009.
- [SZZ11] Roman Schlegel, Kehuan Zhang, Xiaoyong Zhou, Mehool Intwala, Apu Kapadia, and X Wang. “Soundcomber: A Stealthy and Context-Aware Sound Trojan for Smartphones.” In *NDSS*, 2011.
- [TAY12] Chen Tian, Richard Alimi, Yang Richard Yang, and David Zhang. “ShadowStream: Performance Evaluation as a Capability in Production Internet Live Streaming Networks.” In *ACM SIGCOMM*, pp. 347–358, 2012.

- [TB10] Dirk Trossen and Gergely Biczók. “Not Paying the Truck Driver: Differentiated Pricing for the Future Internet.” In *Proceedings of the Re-Architecting the Internet Workshop*, ReARCH ’10, 2010.
- [TEM06] Patrick Traynor, William Enck, Patrick McDaniel, and Thomas La Porta. “Mitigating Attacks on Open Functionality in SMS-capable Cellular Networks.” In *ACM MOBICOM*, pp. 182–193, 2006.
- [The12a] The Wall Street Journal. “Facebook to Target Ads Based on App Usage.”, July 2012. <http://online.wsj.com/article/SB10001424052702304141204577510951953200634.html>.
- [The12b] The Wall Street Journal. “Mobile Ads: Here’s What Works and What Doesn’t.”, Sep 2012. <http://online.wsj.com/article/SB10000872396390444083304578016373342878556.html>.
- [TLO09] Patrick Traynor, Michael Lin, Machigar Ongtang, Vikhyath Rao, Trent Jaeger, Patrick McDaniel, and Thomas La Porta. “On Cellular Botnets: Measuring the Impact of Malicious Devices on a Cellular Network Core.” In *ACM CCS*, pp. 223–234, 2009.
- [TML07] Patrick Traynor, Patrick McDaniel, Thomas La Porta, et al. “On Attack Causality in Internet-Connected Cellular Networks.” In *USENIX Security*, pp. 1–16, 2007.
- [TPL13] Guan hua Tu, Chunyi Peng, Chi yu Li, Xingyu Ma, Hongyi Wang, and Songwu Lu. “Accounting for Roaming Users on Mobile Data Access: Issues and Root Causes.” In *ACM MOBISYS*, 2013.
- [Traa] “Traffic Monitor - RadioOpt GmbH.” <https://play.google.com/store/apps/details?id=com.radioopt.widget>.
- [Trab] TrafficStats. <http://developer.android.com>.
- [VBJ09] Mythili Vutukuru, Hari Balakrishnan, and Kyle Jamieson. “Cross-layer Wireless Bit Rate Adaptation.” In *SIGCOMM*, 2009.
- [VHV12] Balajee Vamanan, Jahangir Hasan, and T.N. Vijaykumar. “Deadline-aware Datacenter Tcp (D2TCP).” In *ACM SIGCOMM*, pp. 115–126, 2012.
- [VLC] VLC Stream & Convert. <http://traveldevel.com/>.
- [VoL] “Voice over LTE.” <http://www.gsma.com/technicalprojects/volte>.
- [VRC08] Luis M. Vaquero, Luis Roderó-Merino, Juan Cáceres, and Maik Lindner. “A Break in the Clouds: Towards a Cloud Definition.” *SIGCOMM Comput. Commun. Rev.*, **39**(1):50–55, December 2008.

- [Wal11] Dan S. Wallach. “Smartphone Security: Trends and Predictions.” In *Secure Application Development*, SecAppDev, 2011.
- [Wik] Wikiversity. “How is Telephony Traffic Modelled?” http://en.wikiversity.org/wiki/Teletraffic_engineering/How_is_Telephony_Traffic_Modelled.
- [Win] Wing FTP. <http://www.wftpserver.com/>.
- [Wir] WireShark/TShark. <http://www.wireshark.org/>.
- [WJP13] Shinae Woo, Eunyoung Jeong, Shinjo Park, Jongmin Lee, Sunghwan Ihm, and Kyoungsoo Park. “Comparison of Caching Strategies in Modern Cellular Backhaul Networks.” In *ACM MOBISYS*, 2013.
- [WQX11] Zhaoguang Wang, Zhiyun Qian, Qiang Xu, Zhuoqing Mao, and Ming Zhang. “An Untold Story of Middleboxes in Cellular Networks.” In *ACM SIGCOMM*, 2011.
- [WVC12] Xiaofei Wang, Athanasios V. Vasilakos, Min Chen, Yunhao Liu, and Ted Taekyoung Kwon. “A Survey of Green Mobile Networks: Opportunities and Challenges.” *Mob. Netw. Appl.*, **17**(1):4–20, February 2012.
- [YDH07] Hung-chih Yang, Ali Dasdan, Ruey-Lung Hsiao, and D. Stott Parker. “Map-Reduce-Merge: Simplified Relational Data Processing on Large Clusters.” In *ACM SIGMOD*, pp. 1029–1040, 2007.
- [YSC11] Jie Yang, Simon Sidhom, Gayathri Chandrasekaran, Tam Vu, Hongbo Liu, Nicolae Cecan, Yingying Chen, Marco Gruteser, and Richard P. Martin. “Detecting Driver Phone Use Leveraging Car Speakers.” In *ACM MOBICOM*, pp. 97–108, 2011.
- [ZB11] Hui Zang and Jean Bolot. “Anonymization of Location Data does not Work: a Large-scale Measurement Study.” In *ACM MOBICOM*, pp. 145–156, 2011.
- [ZCC12] Zengbin Zhang, David Chu, Xiaomeng Chen, and Thomas Moscibroda. “SwordFight: Enabling a New Class of Phone-to-Phone Action Games on Commodity Phones.” In *ACM MOBISYS*, pp. 1–14, 2012.
- [ZXH12] Xinggong Zhang, Yang Xu, Hao Hu, Yong Liu, Zongming Guo, and Yao Wang. “Profiling Skype Video Calls: Rate Control and Video Quality.” In *INFOCOM’12*, March 2012.